

Fast global convergence of gradient methods for high-dimensional statistical recovery

Alekh Agarwal[†]

alekh@eecs.berkeley.edu

Sahand N. Negahban[‡]

sahandn@mit.edu

Martin J. Wainwright^{*,†}

wainwrig@stat.berkeley.edu

Department of Statistics^{*}, and

Department of EECS[†],

University of California, Berkeley
Berkeley, CA 94720-1776

Department of EECS[‡],

Massachusetts Institute of Technology,
32 Vassar Street
Cambridge MA 02139

Abstract

Many statistical M -estimators are based on convex optimization problems formed by the combination of a data-dependent loss function with a norm-based regularizer. We analyze the convergence rates of projected gradient and composite gradient methods for solving such problems, working within a high-dimensional framework that allows the data dimension d to grow with (and possibly exceed) the sample size n . This high-dimensional structure precludes the usual global assumptions—namely, strong convexity and smoothness conditions—that underlie much of classical optimization analysis. We define appropriately restricted versions of these conditions, and show that they are satisfied with high probability for various statistical models. Under these conditions, our theory guarantees that projected gradient descent has a globally geometric rate of convergence up to the *statistical precision* of the model, meaning the typical distance between the true unknown parameter θ^* and an optimal solution $\hat{\theta}$. This result is substantially sharper than previous convergence results, which yielded sublinear convergence, or linear convergence only up to the noise level. Our analysis applies to a wide range of M -estimators and statistical models, including sparse linear regression using Lasso (ℓ_1 -regularized regression); group Lasso for block sparsity; log-linear models with regularization; low-rank matrix recovery using nuclear norm regularization; and matrix decomposition. Overall, our analysis reveals interesting connections between statistical precision and computational efficiency in high-dimensional estimation.

1 Introduction

High-dimensional data sets present challenges that are both statistical and computational in nature. On the statistical side, recent years have witnessed a flurry of results on consistency and rates for various estimators under non-asymptotic high-dimensional scaling, meaning that error bounds are provided for general settings of the sample size n and problem dimension d , allowing for the possibility that $d \gg n$. These results typically involve some assumption regarding the underlying structure of the parameter space, such as sparse vectors, structured covariance matrices, low-rank matrices, or structured regression functions, as well as some regularity conditions on the data-generating process. On the computational side, many estimators for statistical recovery are based on solving convex programs. Examples of such M -estimators include ℓ_1 -regularized quadratic programs (also known as the Lasso) for sparse linear regression (e.g., see the papers [41, 13, 45, 27, 6, 9, 43] and references therein), second-order cone programs (SOCP) for the group Lasso (e.g., [46, 25, 20] and references therein), and semidefinite programming relaxations (SDP) for various problems,

including sparse PCA and low-rank matrix estimation (e.g., [11, 36, 40, 2, 38, 29, 37] and references therein).

Many of these programs are instances of convex conic programs, and so can (in principle) be solved to ϵ -accuracy in polynomial time using interior point methods, and other standard methods from convex programming (e.g., see the books [5, 7]). However, the complexity of such quasi-Newton methods can be prohibitively expensive for the very large-scale problems that arise from high-dimensional data sets. Accordingly, recent years have witnessed a renewed interest in simpler first-order methods, among them the methods of projected gradient descent and mirror descent. Several authors (e.g., [4, 21, 3]) have used variants of Nesterov's accelerated gradient method [32] to obtain algorithms for high-dimensional statistical problems with a sublinear rate of convergence. Note that an optimization algorithm, generating a sequence of iterates $\{\theta^t\}_{t=0}^\infty$, is said to exhibit *sublinear convergence* to an optimum $\hat{\theta}$ if the optimization error $\|\theta^t - \hat{\theta}\|$ decays at the rate $1/t^\kappa$, for some exponent $\kappa > 0$ and norm $\|\cdot\|$. Although this type of convergence is quite slow, it is the best possible with gradient descent-type methods for convex programs under only Lipschitz conditions [31].

It is known that much faster global rates—in particular, a linear or geometric rate—can be achieved if global regularity conditions like strong convexity and smoothness are imposed [31]. An optimization algorithm is said to exhibit *linear or geometric* convergence if the optimization error $\|\theta^t - \hat{\theta}\|$ decays at a rate κ^t , for some contraction coefficient $\kappa \in (0, 1)$. Note that such convergence is exponentially faster than sub-linear convergence. For certain classes of problems involving polyhedral constraints and global smoothness, Tseng and Luo [26] have established geometric convergence. However, a challenging aspect of statistical estimation in high dimensions is that the underlying optimization problems can never be strongly convex in a global sense when $d > n$ (since the $d \times d$ Hessian matrix is rank-deficient), and global smoothness conditions cannot hold when $d/n \rightarrow +\infty$. Some more recent work has exploited structure specific to the optimization problems that arise in statistical settings. For the special case of sparse linear regression with random isotropic designs (also referred to as compressed sensing), some authors have established fast convergence rates in a local sense, meaning guarantees that apply once the iterates are close enough to the optimum [8, 18]. The intuition underlying these results is that once an algorithm identifies the support set of the optimal solution, the problem is then effectively reduced to a lower-dimensional subspace, and thus fast convergence can be guaranteed in a local sense. Also in the setting of compressed sensing, Tropp and Gilbert [42] studied finite convergence of greedy algorithms based on thresholding techniques, and showed linear convergence up to a certain tolerance. For the same class of problems, Garg and Khandekar [17] showed that a thresholded gradient algorithm converges rapidly up to some tolerance. In both of these results, the convergence tolerance is of the order of the noise variance, and hence substantially larger than the true statistical precision of the problem.

The focus of this paper is the convergence rate of two simple gradient-based algorithms for solving optimization problems that underlie regularized M -estimators. For a constrained problem with a differentiable objective function, the projected gradient method generates a sequence of iterates $\{\theta^t\}_{t=0}^\infty$ by taking a step in the negative gradient direction, and then projecting the result onto the constraint set. The composite gradient method of Nesterov [32] is well-suited to solving regularized problems formed by the sum of a differentiable and (potentially) non-differentiable component. The main contribution of this paper is to establish a form of global geometric convergence for these algorithms that holds for a broad class of high-dimensional statistical problems. In order to provide intuition for this guarantee, Figure 1 shows the performance of projected gradient descent

for a Lasso problem (ℓ_1 -constrained least-squares). In panel (a), we have plotted the logarithm of the optimization error, measured in terms of the Euclidean norm $\|\theta^t - \hat{\theta}\|$ between the current iterate θ^t and an optimal solution $\hat{\theta}$, versus the iteration number t . The plot includes three different curves, corresponding to sparse regression problems in dimension $d \in \{5000, 10000, 20000\}$, and a fixed sample size $n = 2500$. Note that all curves are linear (on this logarithmic scale), revealing the geometric convergence predicted by our theory. Such convergence is not predicted by classical optimization theory, since the objective function cannot be strongly convex whenever $n < d$. Moreover, the convergence is geometric even at early iterations, and takes place to a precision far less than the noise level ($\nu^2 = 0.25$ in this example). We also note that the design matrix does not satisfy the restricted isometry property, as assumed in some past work.




Figure 1. Convergence rates of projected gradient descent in application to Lasso programs (ℓ_1 -constrained least-squares). Each panel shows the log optimization error $\log \|\theta^t - \hat{\theta}\|$ versus the iteration number t . Panel (a) shows three curves, corresponding to dimensions $d \in \{5000, 10000, 20000\}$, sparsity $s = \lceil \sqrt{d} \rceil$, and all with the same sample size $n = 2500$. All cases show geometric convergence, but the rate for larger problems becomes progressively slower. (b) For an appropriately rescaled sample size ($\alpha = \frac{n}{s \log d}$), all three convergence rates should be roughly the same, as predicted by the theory.

The results in panel (a) exhibit an interesting property: the convergence rate is *dimension-dependent*, meaning that for a fixed sample size, projected gradient descent converges more slowly for a large problem than a smaller problem—compare the squares for $d = 20000$ to the diamonds for $d = 5000$. This phenomenon reflects the natural intuition that larger problems are, in some sense, “harder” than smaller problems. A notable aspect of our theory is that in addition to guaranteeing geometric convergence, it makes a quantitative prediction regarding the extent to which a larger problem is harder than a smaller one. In particular, our convergence rates suggest that if the sample size n is re-scaled in a certain way according to the dimension d and also other model parameters such as sparsity, then convergence rates should be roughly similar. Panel (b) provides a confirmation of this prediction: when the sample size is rescaled according to our theory (in particular, see Corollary 2 in Section 3.2), then all three curves lie essentially on top of another.

Although high-dimensional optimization problems are typically neither strongly convex nor smooth, this paper shows that it is fruitful to consider suitably restricted notions of strong con-

vexity and smoothness. Our notion of restricted strong convexity (RSC) is related to but slightly different than that introduced in a recent paper by Negahban et al. [28] for establishing statistical consistency. As we discuss in the sequel, bounding the optimization error introduces new challenges not present when analyzing the statistical error. We also introduce a related notion of restricted smoothness (RSM), not needed for proving statistical rates but essential in the setting of optimization. Our analysis consists of two parts. We first show that for optimization problems underlying many regularized M -estimators, appropriately modified notions of restricted strong convexity (RSC) and smoothness (RSM) are sufficient to guarantee global linear convergence of projected gradient descent. Our second contribution is to prove that for the iterates generated by our first-order method, these RSC/RSM assumptions do indeed hold with high probability for a broad class of statistical models, among them sparse linear models, models with group sparsity constraints, and various classes of matrix estimation problems, including matrix completion and matrix decomposition.

An interesting aspect of our results is that the global geometric convergence is not guaranteed to an arbitrary numerical precision, but only to an accuracy related to *statistical precision* of the problem. For a given error norm $\|\cdot\|$, given by the Euclidean or Frobenius norm for most examples in this paper, the statistical precision is given by the mean-squared error $\mathbb{E}[\|\hat{\theta} - \theta^*\|^2]$ between the true parameter θ^* and the estimate $\hat{\theta}$ obtained by solving the optimization problem, where the expectation is taken over randomness in the statistical model. Note that this is very natural from the statistical perspective, since it is the true parameter θ^* itself (as opposed to the solution $\hat{\theta}$ of the M -estimator) that is of primary interest, and our analysis allows us to approach it as close as is statistically possible. Our analysis shows that we can geometrically converge to a parameter θ such that $\|\theta - \theta^*\| = \|\hat{\theta} - \theta^*\| + o(\|\hat{\theta} - \theta^*\|)$, which is the best we can hope for statistically, ignoring lower order terms. Overall, our results reveal an interesting connection between the statistical and computational properties of M -estimators—that is, the properties of the underlying statistical model that make it favorable for estimation also render it more amenable to optimization procedures.

The remainder of this paper is organized as follows. We begin in Section 2 with a precise formulation of the class of convex programs analyzed in this paper, along with background on the notions of a decomposable regularizer, and properties of the loss function. Section 3 is devoted to the statement of our main convergence result, as well as to the development and discussion of its various corollaries for specific statistical models. In Section 4, we provide a number of empirical results that confirm the sharpness of our theoretical predictions. Finally, Section 5 contains the proofs, with more technical aspects of the arguments deferred to the Appendix.

2 Background and problem formulation

In this section, we begin by describing the class of regularized M -estimators to which our analysis applies, as well as the optimization algorithms that we analyze. Finally, we introduce some important notions that underlie our analysis, including the notions of a decomposable regularization, and the properties of restricted strong convexity and smoothness.

2.1 Loss functions, regularization and gradient-based methods

Given a random variable $Z \sim \mathbb{P}$ taking values in some set \mathcal{Z} , let $Z_1^n = \{Z_1, \dots, Z_n\}$ be a collection of n observations. Here the integer n is the *sample size* of the problem. Assuming that \mathbb{P} lies within some indexed family $\{\mathbb{P}_\theta, \theta \in \Omega\}$, the goal is to recover an estimate of the unknown true parameter $\theta^* \in \Omega$ generating the data. Here Ω is some subset of \mathbb{R}^d , and the integer d is known as the *ambient dimension* of the problem. In order to measure the “fit” of any given parameter $\theta \in \Omega$ to a given data set Z_1^n , we introduce a loss function $\mathcal{L}_n : \Omega \times \mathcal{Z}^n \rightarrow \mathbb{R}_+$. By construction, for any given n -sample data set $Z_1^n \in \mathcal{Z}^n$, the loss function assigns a cost $\mathcal{L}_n(\theta; Z_1^n) \geq 0$ to the parameter $\theta \in \Omega$. In many (but not all) applications, the loss function has a separable structure across the data set, meaning that $\mathcal{L}_n(\theta; Z_1^n) = \frac{1}{n} \sum_{i=1}^n \ell(\theta; Z_i)$ where $\ell : \Omega \times \mathcal{Z} \rightarrow \mathbb{R}_+$ is the loss function associated with a single data point.

Of primary interest in this paper are estimation problems that are under-determined, meaning that the number of observations n is smaller than the ambient dimension d . In such settings, without further restrictions on the parameter space Ω , there are various impossibility theorems, asserting that consistent estimates of the unknown parameter θ^* cannot be obtained. For this reason, it is necessary to assume that the unknown parameter θ^* either lies within a smaller subset of Ω , or is well-approximated by some member of such a subset. In order to incorporate these types of structural constraints, we introduce a *regularizer* $\mathcal{R} : \Omega \rightarrow \mathbb{R}_+$ over the parameter space. With these ingredients, the analysis of this paper applies to the *constrained M-estimator*

$$\hat{\theta}_\rho \in \arg \min_{\mathcal{R}(\theta) \leq \rho} \{\mathcal{L}_n(\theta; Z_1^n)\}, \quad (1)$$

where $\rho > 0$ is a user-defined radius, as well as to the *regularized M-estimator*

$$\hat{\theta}_{\lambda_n} \in \arg \min_{\mathcal{R}(\theta) \leq \bar{\rho}} \underbrace{\{\mathcal{L}_n(\theta; Z_1^n) + \lambda_n \mathcal{R}(\theta)\}}_{\phi_n(\theta)} \quad (2)$$

where the regularization weight $\lambda_n > 0$ is user-defined. Note that the radii ρ and $\bar{\rho}$ may be different in general. Throughout this paper, we impose the following two conditions:

- (a) for any data set Z_1^n , the function $\mathcal{L}_n(\cdot; Z_1^n)$ is convex and differentiable over Ω , and
- (b) the regularizer \mathcal{R} is a norm.

These conditions ensure that the overall problem is convex, so that by Lagrangian duality, the optimization problems (1) and (2) are equivalent. However, as our analysis will show, solving one or the other can be computationally more preferable depending upon the assumptions made. Some remarks on notation: when the radius ρ or the regularization parameter λ_n is clear from the context, we will drop the subscript on $\hat{\theta}$ to ease the notation. Similarly, we frequently adopt the shorthand $\mathcal{L}_n(\theta)$, with the dependence of the loss function on the data being implicitly understood. Procedures based on optimization problems of either form are known as *M-estimators* in the statistics literature.

The focus of this paper is on two simple algorithms for solving the above optimization problems. The method of *projected gradient descent* applies naturally to the constrained problem (1), whereas the *composite gradient descent* method due to Nesterov [32] is suitable for solving the regularized

problem (2). Each routine generates a sequence $\{\theta^t\}_{t=0}^\infty$ of iterates by first initializing to some parameter $\theta^0 \in \Omega$, and then applying the recursive update

$$\theta^{t+1} = \arg \min_{\theta \in \mathbb{B}_\mathcal{R}(\rho)} \{\mathcal{L}_n(\theta^t) + \langle \nabla \mathcal{L}_n(\theta^t), \theta - \theta^t \rangle + \frac{\gamma_u}{2} \|\theta - \theta^t\|^2\}, \quad \text{for } t = 0, 1, 2, \dots, \quad (3)$$

in the case of projected gradient descent, or the update

$$\theta^{t+1} = \arg \min_{\theta \in \mathbb{B}_\mathcal{R}(\bar{\rho})} \{\mathcal{L}_n(\theta^t) + \langle \nabla \mathcal{L}_n(\theta^t), \theta - \theta^t \rangle + \frac{\gamma_u}{2} \|\theta - \theta^t\|^2 + \lambda_n \mathcal{R}(\theta)\}, \quad \text{for } t = 0, 1, 2, \dots, \quad (4)$$

for the composite gradient method. Note that the only difference between the two updates is the addition of the regularization term in the objective. These updates have a natural intuition: the next iterate θ^{t+1} is obtained by constrained minimization of a first-order approximation to the loss function, combined with a smoothing term that controls how far one moves from the current iterate in terms of Euclidean norm. Moreover, it is easily seen that the update (3) is equivalent to

$$\theta^{t+1} = \Pi \left(\theta^t - \frac{1}{\gamma_u} \nabla \mathcal{L}_n(\theta^t) \right), \quad (5)$$

where $\Pi \equiv \Pi_{\mathbb{B}_\mathcal{R}(\rho)}$ denotes Euclidean projection onto the ball $\mathbb{B}_\mathcal{R}(\rho) = \{\theta \in \Omega \mid \mathcal{R}(\theta) \leq \rho\}$ of radius ρ . In this formulation, we see that the algorithm takes a step in the negative gradient direction, using the quantity $1/\gamma_u$ as stepsize parameter, and then projects the resulting vector onto the constraint set. The update (4) takes an analogous form, however, the projection will depend on both λ_n and γ_u . As will be illustrated in the examples to follow, for many problems, the updates (3) and (4), or equivalently (5), have a very simple solution. For instance, in the case of ℓ_1 -regularization, it can be obtained by an appropriate form of the soft-thresholding operator.

2.2 Restricted strong convexity and smoothness

In this section, we define the conditions on the loss function and regularizer that underlie our analysis. Global smoothness and strong convexity assumptions play an important role in the classical analysis of optimization algorithms [5, 7, 31]. In application to a differentiable loss function \mathcal{L}_n , both of these properties are defined in terms of a first-order Taylor series expansion around a vector θ' in the direction of θ —namely, the quantity

$$\mathcal{T}_\mathcal{L}(\theta; \theta') := \mathcal{L}_n(\theta) - \mathcal{L}_n(\theta') - \langle \nabla \mathcal{L}_n(\theta'), \theta - \theta' \rangle. \quad (6)$$

By the assumed convexity of \mathcal{L}_n , this error is always non-negative, and global strong convexity is equivalent to imposing a stronger condition, namely that for some parameter $\gamma_\ell > 0$, the first-order Taylor error $\mathcal{T}_\mathcal{L}(\theta; \theta')$ is lower bounded by a quadratic term $\frac{\gamma_\ell}{2} \|\theta - \theta'\|^2$ for all $\theta, \theta' \in \Omega$. Global smoothness is defined in a similar way, by imposing a quadratic upper bound on the Taylor error. It is known that under global smoothness and strong convexity assumptions, the method of projected gradient descent (3) enjoys a *globally geometric convergence rate*, meaning that there is some $\kappa \in (0, 1)$ such that¹

$$\|\theta^t - \hat{\theta}\|^2 \lesssim \kappa^t \|\theta^0 - \hat{\theta}\|^2 \quad \text{for all iterations } t = 0, 1, 2, \dots. \quad (7)$$

¹In this statement (and throughout the paper), we use \lesssim to mean an inequality that holds with some universal constant c , independent of the problem parameters.

We refer the reader to Bertsekas [5, Prop. 1.2.3, p. 145], or Nesterov [31, Thm. 2.2.8, p. 88] for such results on projected gradient descent, and to Nesterov [32] for composite gradient descent.

Unfortunately, in the high-dimensional setting ($d > n$), it is usually impossible to guarantee strong convexity of the problem (1) in a global sense. For instance, when the data is drawn i.i.d., the loss function consists of a sum of n terms. If the loss is twice differentiable, the resulting $d \times d$ Hessian matrix $\nabla^2 \mathcal{L}(\theta; Z_1^n)$ is often a sum of n matrices each with rank one, so that the Hessian is rank-degenerate when $n < d$. However, as we show in this paper, in order to obtain fast convergence rates for the optimization method (3), it is sufficient that (a) the objective is strongly convex and smooth in a restricted set of directions, and (b) the algorithm approaches the optimum $\hat{\theta}$ only along these directions. Let us now formalize these ideas.

Definition 1 (Restricted strong convexity (RSC)). *The loss function \mathcal{L}_n satisfies restricted strong convexity with respect to \mathcal{R} and with parameters $(\gamma_\ell, \tau_\ell(\mathcal{L}_n))$ over the set Ω' if*

$$\mathcal{T}_{\mathcal{L}}(\theta; \theta') \geq \frac{\gamma_\ell}{2} \|\theta - \theta'\|^2 - \tau_\ell(\mathcal{L}_n) \mathcal{R}^2(\theta - \theta') \quad \text{for all } \theta, \theta' \in \Omega'. \quad (8)$$

We refer to the quantity γ_ℓ as the *(lower) curvature parameter*, and to the quantity τ_ℓ as the *tolerance parameter*. The set Ω' corresponds to a suitably chosen subset of the space Ω of all possible parameters.

In order to gain intuition for this definition, first suppose that the condition (8) holds with tolerance parameter $\tau_\ell = 0$. In this case, the regularizer plays no role in the definition, and condition (8) is equivalent to the usual definition of strong convexity on the optimization set Ω . As discussed previously, this type of global strong convexity typically *fails* to hold for high-dimensional inference problems. In contrast, when tolerance parameter τ_ℓ is strictly positive, the condition (8) is much milder, in that it only applies to a *limited set* of vectors. For a given pair $\theta \neq \theta'$, consider the inequality

$$\frac{\mathcal{R}^2(\theta - \theta')}{\|\theta - \theta'\|^2} < \frac{\gamma_\ell}{2 \tau_\ell(\mathcal{L}_n)}. \quad (9)$$

If this inequality is violated, then the right-hand side of the bound (8) is non-positive, in which case the RSC constraint (8) is vacuous. Thus, restricted strong convexity imposes a non-trivial constraint only on pairs $\theta \neq \theta'$ for which the inequality (8) holds, and a central part of our analysis will be to prove that, for the sequence of iterates generated by projected gradient descent, the optimization error $\hat{\Delta}^t := \theta^t - \hat{\theta}$ satisfies a constraint of the form (9). We note that since the regularizer \mathcal{R} is convex, strong convexity of the loss function \mathcal{L}_n also implies the strong convexity of the regularized loss ϕ_n as well.

For the least-squares loss, the RSC definition depends purely on the direction (and not the magnitude) of the difference vector $\theta - \theta'$. For other types of loss functions—such as those arising in generalized linear models—it is essential to localize the RSC definition, requiring that it holds only for pairs for which the norm $\|\theta - \theta'\|_2$ is not too large. We refer the reader to Section 2.4.1 for further discussion of this issue.

Finally, as pointed out by a reviewer, our restricted version of strong convexity can be seen as an instance of the general theory of paraconvexity (e.g., [33]); however, we are not aware of convergence rates for minimizing general paraconvex functions.

We also specify an analogous notion of restricted smoothness:

Definition 2 (Restricted smoothness (RSM)). We say the loss function \mathcal{L}_n satisfies restricted smoothness with respect to \mathcal{R} and with parameters $(\gamma_u, \tau_u(\mathcal{L}_n))$ over the set Ω' if

$$\mathcal{T}_{\mathcal{L}}(\theta; \theta') \leq \frac{\gamma_u}{2} \|\theta - \theta'\|^2 + \tau_u(\mathcal{L}_n) \mathcal{R}^2(\theta - \theta') \quad \text{for all } \theta, \theta' \in \Omega'. \quad (10)$$

As with our definition of restricted strong convexity, the additional tolerance $\tau_u(\mathcal{L}_n)$ is not present in analogous smoothness conditions in the optimization literature, but it is essential in our set-up.

2.3 Decomposable regularizers

In past work on the statistical properties of regularization, the notion of a decomposable regularizer has been shown to be useful [28]. Although the focus of this paper is a rather different set of questions—namely, optimization as opposed to statistics—decomposability also plays an important role here. Decomposability is defined with respect to a pair of subspaces defined with respect to the parameter space $\Omega \subseteq \mathbb{R}^d$. The set \mathcal{M} is known as the *model subspace*, whereas the set $\overline{\mathcal{M}}^\perp$, referred to as the *perturbation subspace*, captures deviations away from the model subspace.

Definition 3. Given a subspace pair $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$ such that $\mathcal{M} \subseteq \overline{\mathcal{M}}$, we say that a norm \mathcal{R} is $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$ -decomposable if

$$\mathcal{R}(\alpha + \beta) = \mathcal{R}(\alpha) + \mathcal{R}(\beta) \quad \text{for all } \alpha \in \mathcal{M} \text{ and } \beta \in \overline{\mathcal{M}}^\perp. \quad (11)$$

To gain some intuition for this definition, note that by triangle inequality, we always have the bound $\mathcal{R}(\alpha + \beta) \leq \mathcal{R}(\alpha) + \mathcal{R}(\beta)$. For a decomposable regularizer, this inequality always holds with equality. Thus, given a fixed vector $\alpha \in \mathcal{M}$, the key property of any decomposable regularizer is that it affords the *maximum penalization* of any deviation $\beta \in \overline{\mathcal{M}}^\perp$.

For a given error norm $\|\cdot\|$, its interaction with the regularizer \mathcal{R} plays an important role in our results. In particular, we have the following:

Definition 4 (Subspace compatibility). Given the regularizer $\mathcal{R}(\cdot)$ and a norm $\|\cdot\|$, the associated subspace compatibility is given by

$$\Psi(\overline{\mathcal{M}}) := \sup_{\theta \in \overline{\mathcal{M}} \setminus \{0\}} \frac{\mathcal{R}(\theta)}{\|\theta\|} \quad \text{when } \overline{\mathcal{M}} \neq \{0\}, \quad \text{and } \Psi(\{0\}) := 0. \quad (12)$$

The quantity $\Psi(\overline{\mathcal{M}})$ corresponds to the Lipschitz constant of the norm \mathcal{R} with respect to $\|\cdot\|$, when restricted to the subspace $\overline{\mathcal{M}}$.

2.4 Some illustrative examples

We now describe some particular examples of M -estimators with decomposable regularizers, and discuss the form of the projected gradient updates as well as RSC/RSM conditions. We cover two main families of examples: log-linear models with sparsity constraints and ℓ_1 -regularization (Section 2.4.1), and matrix regression problems with nuclear norm regularization (Section 2.4.2).

2.4.1 Sparse log-linear models and ℓ_1 -regularization

Suppose that each sample Z_i consists of a scalar-vector pair $(y_i, x_i) \in \mathbb{R} \times \mathbb{R}^d$, corresponding to the scalar response $y_i \in \mathcal{Y}$ associated with a vector of predictors $x_i \in \mathbb{R}^d$. A log-linear model with canonical link function assumes that the response y_i is linked to the covariate vector x_i via a conditional distribution of the form $\mathbb{P}(y_i \mid x_i; \theta^*, \sigma) \propto \exp\left\{\frac{y_i \langle \theta^*, x_i \rangle - \Phi(\langle \theta^*, x_i \rangle)}{c(\sigma)}\right\}$, where $c(\sigma)$ is a known quantity, $\Phi(\cdot)$ is the log-partition function to normalize the density, and $\theta^* \in \mathbb{R}^d$ is an unknown regression vector. In many applications, the regression vector θ^* is relatively sparse, so that it is natural to impose an ℓ_1 -constraint. Computing the maximum likelihood estimate subject to such a constraint involves solving the convex program²

$$\hat{\theta} \in \arg \min_{\theta \in \Omega} \underbrace{\left\{ \frac{1}{n} \sum_{i=1}^n \{\Phi(\langle \theta, x_i \rangle) - y_i \langle \theta, x_i \rangle\} \right\}}_{\mathcal{L}_n(\theta; Z_1^n)} \quad \text{such that } \|\theta\|_1 \leq \rho, \quad (13)$$

with $x_i \in \mathbb{R}^d$ as its i^{th} row. We refer to this estimator as the log-linear Lasso; it is a special case of the M -estimator (1), with the loss function $\mathcal{L}_n(\theta; Z_1^n) = \frac{1}{n} \sum_{i=1}^n \{\Phi(\langle \theta, x_i \rangle) - y_i \langle \theta, x_i \rangle\}$ and the regularizer $\mathcal{R}(\theta) = \|\theta\|_1 = \sum_{j=1}^d |\theta_j|$.

Ordinary linear regression is the special case of the log-linear setting with $\Phi(t) = t^2/2$ and $\Omega = \mathbb{R}^d$, and in this case, the estimator (13) corresponds to ordinary least-squares version of Lasso [13, 41]. Other forms of log-linear Lasso that are of interest include logistic regression, Poisson regression, and multinomial regression.

Projected gradient updates: Computing the gradient of the log-linear loss from equation (13) is straightforward: we have $\nabla \mathcal{L}_n(\theta) = \frac{1}{n} \sum_{i=1}^n x_i \{\Phi'(\langle \theta, x_i \rangle) - y_i\}$, and the update (5) corresponds to the Euclidean projection of the vector $\theta^t - \frac{1}{\gamma_u} \nabla \mathcal{L}_n(\theta^t)$ onto the ℓ_1 -ball of radius ρ . It is well-known that this projection can be characterized in terms of soft-thresholding, and that the projected update (5) can be computed easily. We refer the reader to Duchi et al. [14] for an efficient implementation requiring $\mathcal{O}(d)$ operations.

Composite gradient updates: The composite gradient update for this problem amounts to solving

$$\theta^{t+1} = \arg \min_{\|\theta\|_1 \leq \bar{\rho}} \left\{ \langle \theta, \nabla \mathcal{L}_n(\theta) \rangle + \frac{\gamma_u}{2} \|\theta - \theta^t\|_2^2 + \lambda_n \|\theta\|_1 \right\}.$$

The update can be computed by two soft-thresholding operations. The first step is soft thresholding the vector $\theta^t - \frac{1}{\gamma_u} \nabla \mathcal{L}_n(\theta^t)$ at a level λ_n . If the resulting vector has ℓ_1 -norm greater than $\bar{\rho}$, then we project on to the ℓ_1 -ball just like before. Overall, the complexity of the update is still $\mathcal{O}(d)$ as before.

Decomposability of ℓ_1 -norm: We now illustrate how the ℓ_1 -norm is decomposable with respect to appropriately chosen subspaces. For any subset $S \subseteq \{1, 2, \dots, d\}$, consider the subspace

$$\mathcal{M}(S) := \{\alpha \in \mathbb{R}^d \mid \alpha_j = 0 \text{ for all } j \notin S\}, \quad (14)$$

²The link function Φ is convex since it is the log-partition function of a canonical exponential family.

corresponding to all vectors supported only on S . Defining $\overline{\mathcal{M}}(S) = \mathcal{M}(S)$, its orthogonal complement (with respect to the usual Euclidean inner product) is given by

$$\overline{\mathcal{M}}^\perp(S) = \mathcal{M}^\perp(S) = \{\beta \in \mathbb{R}^d \mid \beta_j = 0 \text{ for all } j \in S\}. \quad (15)$$

To establish the decomposability of the ℓ_1 -norm with respect to the pair $(\mathcal{M}(S), \overline{\mathcal{M}}^\perp(S))$, note that any $\alpha \in \mathcal{M}(S)$ can be written in the partitioned form $\alpha = (\alpha_S, 0_{S^c})$, where $\alpha_S \in \mathbb{R}^s$ and $0_{S^c} \in \mathbb{R}^{d-s}$ is a vector of zeros. Similarly, any vector $\beta \in \overline{\mathcal{M}}^\perp(S)$ has the partitioned representation $(0_S, \beta_{S^c})$. With these representations, we have the decomposition

$$\|\alpha + \beta\|_1 = \|(\alpha_S, 0) + (0, \beta_{S^c})\|_1 = \|\alpha\|_1 + \|\beta\|_1.$$

Consequently, for any subset S , the ℓ_1 -norm is decomposable with respect to the pairs $(\mathcal{M}(S), \mathcal{M}^\perp(S))$.

In analogy to the ℓ_1 -norm, various types of group-sparse norms are also decomposable with respect to non-trivial subspace pairs. We refer the reader to the paper [28] for further discussion and examples of such decomposable norms.

RSC/RSM conditions: A calculation using the mean-value theorem shows that for the loss function (13), the error in the first-order Taylor series, as previously defined in equation (6), can be written as

$$\mathcal{T}_{\mathcal{L}}(\theta; \theta') = \frac{1}{n} \sum_{i=1}^n \Phi''(\langle \theta_t, x_i \rangle) (\langle x_i, \theta - \theta' \rangle)^2$$

where $\theta_t = t\theta + (1-t)\theta'$ for some $t \in [0, 1]$. When $n < d$, then we can always find pairs $\theta \neq \theta'$ such that $\langle x_i, \theta - \theta' \rangle = 0$ for all $i = 1, 2, \dots, n$, showing that the objective function can never be strongly convex. On the other hand, restricted strong convexity for log-linear models requires only that there exist positive numbers $(\gamma_\ell, \tau_\ell(\mathcal{L}_n))$ such that

$$\frac{1}{n} \sum_{i=1}^n \Phi''(\langle \theta_t, x_i \rangle) (\langle x_i, \theta - \theta' \rangle)^2 \geq \frac{\gamma_\ell}{2} \|\theta - \theta'\|^2 - \tau_\ell(\mathcal{L}_n) \mathcal{R}^2(\theta - \theta') \quad \text{for all } \theta, \theta' \in \Omega', \quad (16)$$

where $\Omega' := \Omega \cap \mathbb{B}_2(R)$ is the intersection of the parameter space Ω with a Euclidean ball of some fixed radius R around zero. This restriction is essential because for many generalized linear models, the Hessian function Φ'' approaches zero as its argument diverges. For instance, for the logistic function $\Phi(t) = \log(1 + \exp(t))$, we have $\Phi''(t) = \exp(t)/[1 + \exp(t)]^2$, which tends to zero as $t \rightarrow +\infty$. Restricted smoothness imposes an analogous upper bound on the Taylor error. For a broad class of log-linear models, such bounds hold with tolerance $\tau_\ell(\mathcal{L}_n)$ and $\tau_u(\mathcal{L}_n)$ of the order $\sqrt{\frac{\log d}{n}}$. Further details on such results are provided in the corollaries to follow our main theorem. A detailed discussion of RSC for exponential families in statistical problems can be found in the paper [28].

In order to ensure RSC/RSM conditions on the iterates θ^t of the updates (3) or (4), we also need to ensure that $\theta^t \in \Omega'$. This can be done by defining $\mathcal{L}'_n = \mathcal{L}_n + \mathbb{I}_{\Omega'}(\theta)$, where $\mathbb{I}_{\Omega'}(\theta)$ is zero when $\theta \in \Omega'$ and ∞ otherwise. This is equivalent to projection on the intersection of ℓ_1 -ball with Ω' in the updates (3) and (4) and can be done efficiently with Dykstra's algorithm [15], for instance, as long as the individual projections are efficient.

In the special case of linear regression, we have $\Phi''(t) = 1$ for all $t \in \mathbb{R}$, so that the lower bound (16) involves only the Gram matrix $X^T X/n$. (Here $X \in \mathbb{R}^{n \times d}$ is the usual design matrix, with $x_i \in \mathbb{R}^d$ as its i^{th} row.) For linear regression and ℓ_1 -regularization, the RSC condition is equivalent to the lower bound

$$\frac{\|X(\theta - \theta')\|_2^2}{n} \geq \frac{\gamma_\ell}{2} \|\theta - \theta'\|_2^2 - \tau_\ell(\mathcal{L}_n) \|\theta - \theta'\|_1^2 \quad \text{for all } \theta, \theta' \in \Omega. \quad (17)$$

Such a condition corresponds to a variant of the restricted eigenvalue (RE) conditions that have been studied in the literature [6, 43]. Such RE conditions are significantly milder than the restricted isometry property; we refer the reader to van de Geer and Bühlmann [43] for an in-depth comparison of different RE conditions. From past work, the condition (17) is satisfied with high probability for a broad classes of anisotropic random design matrices [34, 39], and parts of our analysis make use of this fact.

2.4.2 Matrices and nuclear norm regularization

We now discuss a general class of matrix regression problems that falls within our framework. Consider the space of $d_1 \times d_2$ matrices endowed with the trace inner product $\langle\langle A, B \rangle\rangle := \text{trace}(A^T B)$. In order to ease notation, we define $d := \min\{d_1, d_2\}$. Let $\Theta^* \in \mathbb{R}^{d_1 \times d_2}$ be an unknown matrix and suppose that for $i = 1, 2, \dots, n$, we observe a scalar-matrix pair $Z_i = (y_i, X_i) \in \mathbb{R} \times \mathbb{R}^{d_1 \times d_2}$ linked to Θ^* via the linear model

$$y_i = \langle\langle X_i, \Theta^* \rangle\rangle + w_i, \quad \text{for } i = 1, 2, \dots, n, \quad (18)$$

where w_i is an additive observation noise. In many contexts, it is natural to assume that Θ^* is exactly low-rank, or approximately so, meaning that it is well-approximated by a matrix of low rank. In such settings, a number of authors (e.g., [16, 38, 29]) have studied the M -estimator

$$\widehat{\Theta} \in \arg \min_{\Theta \in \mathbb{R}^{d_1 \times d_2}} \left\{ \frac{1}{2n} \sum_{i=1}^n (y_i - \langle\langle X_i, \Theta \rangle\rangle)^2 \right\} \quad \text{such that } \|\Theta\|_1 \leq \rho, \quad (19)$$

or the corresponding regularized version. Here the *nuclear or trace norm* is given by $\|\Theta\|_1 := \sum_{j=1}^d \sigma_j(\Theta)$, corresponding to the sum of the singular values. This optimization problem is an instance of a semidefinite program. As discussed in more detail in Section 3.3, there are various applications in which this estimator and variants thereof have proven useful.

Form of projected gradient descent: For the M -estimator (19), the projected gradient updates take a very simple form—namely

$$\Theta^{t+1} = \Pi \left(\Theta^t - \frac{1}{\gamma_u} \frac{\sum_{i=1}^n (y_i - \langle\langle X_i, \Theta^t \rangle\rangle) X_i}{n} \right), \quad (20)$$

where Π denotes Euclidean projection onto the nuclear norm ball $\mathbb{B}_1(\rho) = \{\Theta \in \mathbb{R}^{d_1 \times d_2} \mid \|\Theta\|_1 \leq \rho\}$. This nuclear norm projection can be obtained by first computing the singular value decomposition (SVD), and then projecting the vector of singular values onto the ℓ_1 -ball. The latter step can be achieved by the fast projection algorithms discussed earlier, and there are various methods for fast computation of SVDs. The composite gradient update also has a simple form, requiring at most two singular value thresholding operations as was the case for linear regression.

Decomposability of nuclear norm: We now define matrix subspaces for which the nuclear norm is decomposable. Given a target matrix Θ^* —that is, a quantity to be estimated—consider its singular value decomposition $\Theta^* = UDV^T$, where the matrix $D \in \mathbb{R}^{d \times d}$ is diagonal, with the ordered singular values of Θ^* along its diagonal, and $d := \min\{d_1, d_2\}$. For an integer $r \in \{1, 2, \dots, d\}$, let $U^r \in \mathbb{R}^{d \times r}$ denote the matrix formed by the top r left singular vectors of Θ^* in its columns, and we define the matrix V^r in a similar fashion. Using col to denote the column span of a matrix, we then define the subspaces³

$$\mathcal{M}(U^r, V^r) := \{\Theta \in \mathbb{R}^{d_1 \times d_2} \mid \text{col}(\Theta^T) \subseteq \text{col}(V^r), \text{col}(\Theta) \subseteq \text{col}(U^r)\}, \quad \text{and} \quad (21a)$$

$$\overline{\mathcal{M}}^\perp(U^r, V^r) := \{\Theta \in \mathbb{R}^{d_1 \times d_2} \mid \text{col}(\Theta^T) \subseteq (\text{col}(V^r))^\perp, \text{col}(\Theta) \subseteq (\text{col}(U^r))^\perp\}. \quad (21b)$$

Finally, let us verify the decomposability of the nuclear norm. By construction, any pair of matrices $\Theta \in \mathcal{M}(U^r, V^r)$ and $\Gamma \in \overline{\mathcal{M}}^\perp(U^r, V^r)$ have orthogonal row and column spaces, which implies the required decomposability condition—namely $\|\Theta + \Gamma\|_1 = \|\Theta\|_1 + \|\Gamma\|_1$.

In some special cases such as matrix completion or matrix decomposition that we describe in the sequel, Ω' will involve an additional bound on the entries of Θ^* as well as the iterates Θ^t to establish RSC/RSM conditions. This can be done by augmenting the loss with an indicator of the constraint and using cyclic projections for computing the updates as mentioned earlier in Example 2.4.1.

3 Main results and some consequences

We are now equipped to state the two main results of our paper, and discuss some of their consequences. We illustrate its application to several statistical models, including sparse regression (Section 3.2), matrix estimation with rank constraints (Section 3.3), and matrix decomposition problems (Section 3.4).

3.1 Geometric convergence

Recall that the projected gradient algorithm (3) is well-suited to solving an M -estimation problem in its constrained form, whereas the composite gradient algorithm (4) is appropriate for a regularized problem. Accordingly, let $\widehat{\theta}$ be any optimal solution to the constrained problem (1), or the regularized problem (2), and let $\{\theta^t\}_{t=0}^\infty$ be a sequence of iterates generated by the projected gradient updates (3), or the composite gradient updates (4), respectively. Of primary interest to us in this paper are bounds on the *optimization error*, which can be measured either in terms of the error vector $\widehat{\Delta}^t := \theta^t - \widehat{\theta}$, or the difference between the cost of θ^t and the optimal cost defined by $\widehat{\theta}$. In this section, we state two main results—Theorems 1 and 2—corresponding to the constrained and regularized cases respectively. In addition to the optimization error previously discussed, both of these results involve the *statistical error* $\Delta^* := \widehat{\theta} - \theta^*$ between the optimum $\widehat{\theta}$ and the nominal parameter θ^* . At a high level, these results guarantee that under the RSC/RSM conditions, the optimization error shrinks geometrically, with a contraction coefficient that depends on the loss function \mathcal{L}_n via the parameters $(\gamma_\ell, \tau_\ell(\mathcal{L}_n))$ and $(\gamma_u, \tau_u(\mathcal{L}_n))$. An interesting feature is that the contraction occurs only up to a certain tolerance ϵ^2 depending on these same parameters, and the statistical error. However, as we discuss, for many statistical problems of interest, we can

³ Note that the model space $\mathcal{M}(U^r, V^r)$ is *not equal* to $\overline{\mathcal{M}}(U^r, V^r)$. Nonetheless, as required by Definition 3, we do have the inclusion $\mathcal{M}(U^r, V^r) \subseteq \overline{\mathcal{M}}(U^r, V^r)$.

show that this tolerance ϵ^2 is of a lower order than the intrinsic statistical error, and hence can be neglected from the statistical point of view. Consequently, our theory gives an explicit upper bound on the number of iterations required to solve an M -estimation problem up to the statistical precision.

Convergence rates for projected gradient: We now provide the notation necessary for a precise statement of this claim. Our main result actually involves a family of upper bounds on the optimization error, one for each pair $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$ of \mathcal{R} -decomposable subspaces (see Definition 3). As will be clarified in the sequel, this subspace choice can be optimized for different models so as to obtain the tightest possible bounds. For a given pair $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$ such that $16\Psi^2(\overline{\mathcal{M}})\tau_u(\mathcal{L}_n) < \gamma_u$, let us define the *contraction coefficient*

$$\kappa(\mathcal{L}_n; \overline{\mathcal{M}}) := \left\{ 1 - \frac{\gamma_\ell}{\gamma_u} + \frac{16\Psi^2(\overline{\mathcal{M}})(\tau_u(\mathcal{L}_n) + \tau_\ell(\mathcal{L}_n))}{\gamma_u} \right\} \left\{ 1 - \frac{16\Psi^2(\overline{\mathcal{M}})\tau_u(\mathcal{L}_n)}{\gamma_u} \right\}^{-1}. \quad (22)$$

In addition, we define the *tolerance parameter*

$$\epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}}) := \frac{32(\tau_u(\mathcal{L}_n) + \tau_\ell(\mathcal{L}_n)) (2\mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)) + \Psi(\overline{\mathcal{M}})\|\Delta^*\| + 2\mathcal{R}(\Delta^*))^2}{\gamma_u}, \quad (23)$$

where $\Delta^* = \widehat{\theta} - \theta^*$ is the statistical error, and $\Pi_{\mathcal{M}^\perp}(\theta^*)$ denotes the Euclidean projection of θ^* onto the subspace \mathcal{M}^\perp .

In terms of these two ingredients, we now state our first main result:

Theorem 1. *Suppose that the loss function \mathcal{L}_n satisfies the RSC/RSM condition with parameters $(\gamma_\ell, \tau_\ell(\mathcal{L}_n))$ and $(\gamma_u, \tau_u(\mathcal{L}_n))$ respectively. Let $(\mathcal{M}, \overline{\mathcal{M}})$ be any \mathcal{R} -decomposable pair of subspaces such that $\mathcal{M} \subseteq \overline{\mathcal{M}}$ and $0 < \kappa \equiv \kappa(\mathcal{L}_n, \overline{\mathcal{M}}) < 1$. Then for any optimum $\widehat{\theta}$ of the problem (1) for which the constraint is active, we have*

$$\|\theta^{t+1} - \widehat{\theta}\|^2 \leq \kappa^t \|\theta^0 - \widehat{\theta}\|^2 + \frac{\epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}})}{1 - \kappa} \quad \text{for all iterations } t = 0, 1, 2, \dots \quad (24)$$

Remarks: Theorem 1 actually provides a family of upper bounds, one for each \mathcal{R} -decomposable pair $(\mathcal{M}, \overline{\mathcal{M}})$ such that $0 < \kappa \equiv \kappa(\mathcal{L}_n, \overline{\mathcal{M}}) < 1$. This condition is always satisfied by setting $\overline{\mathcal{M}}$ equal to the trivial subspace $\{0\}$: indeed, by definition (12) of the subspace compatibility, we have $\Psi(\overline{\mathcal{M}}) = 0$, and hence $\kappa(\mathcal{L}_n; \{0\}) = (1 - \frac{\gamma_\ell}{\gamma_u}) < 1$. Although this choice of $\overline{\mathcal{M}}$ minimizes the contraction coefficient, it will lead⁴ to a very large tolerance parameter $\epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}})$. A more typical application of Theorem 1 involves non-trivial choices of the subspace $\overline{\mathcal{M}}$.

The bound (24) guarantees that the optimization error decreases geometrically, with contraction factor $\kappa \in (0, 1)$, up to a certain tolerance proportional to $\epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}})$, as illustrated in Figure 2(a). The contraction factor κ approaches the $1 - \gamma_\ell/\gamma_u$ as the number of samples grows. The appearance of the ratio γ_ℓ/γ_u is natural since it measures the conditioning of the objective function; more specifically, it is essentially a restricted condition number of the Hessian matrix. On

⁴Indeed, the setting $\mathcal{M}^\perp = \mathbb{R}^d$ means that the term $\mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)) = \mathcal{R}(\theta^*)$ appears in the tolerance; this quantity is far larger than statistical precision.

the other hand, the tolerance parameter ϵ depends on the choice of decomposable subspaces, the parameters of the RSC/RSM conditions, and the statistical error $\Delta^* = \hat{\theta} - \theta^*$ (see equation (23)). In the corollaries of Theorem 1 to follow, we show that the subspaces can often be chosen such that $\epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}}) = o(\|\hat{\theta} - \theta^*\|^2)$. Consequently, the bound (24) guarantees geometric convergence up to a tolerance *smaller than statistical precision*, as illustrated in Figure 2(b). This is sensible, since in statistical settings, there is no point to optimizing beyond the statistical precision.




Figure 2. (a) Generic illustration of Theorem 1. The optimization error $\hat{\Delta}^t = \theta^t - \hat{\theta}$ is guaranteed to decrease geometrically with coefficient $\kappa \in (0, 1)$, up to the tolerance $\epsilon^2 = \epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}})$, represented by the circle. (b) Relation between the optimization tolerance $\epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}})$ (solid circle) and the statistical precision $\|\Delta^*\| = \|\theta^* - \hat{\theta}\|$ (dotted circle). In many settings, we have $\epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}}) \ll \|\Delta^*\|^2$, so that convergence is guaranteed up to a tolerance lower than statistical precision.

The result of Theorem 1 takes a simpler form when there is a subspace \mathcal{M} that includes θ^* , and the \mathcal{R} -ball radius is chosen such that $\rho \leq \mathcal{R}(\theta^*)$. In this case, by appropriately controlling the error term, we can establish that it is of lower order than the statistical precision —namely, the squared difference $\|\hat{\theta} - \theta^*\|^2$ between an optimal solution $\hat{\theta}$ to the convex program (1), and the unknown parameter θ^* .

Corollary 1. *In addition to the conditions of Theorem 1, suppose that $\theta^* \in \mathcal{M}$ and $\rho \leq \mathcal{R}(\theta^*)$. Then as long as $\Psi^2(\overline{\mathcal{M}})(\tau_u(\mathcal{L}_n) + \tau_\ell(\mathcal{L}_n)) = o(1)$, we have*

$$\|\theta^{t+1} - \hat{\theta}\|^2 \leq \kappa^t \|\theta^0 - \hat{\theta}\|^2 + o(\|\hat{\theta} - \theta^*\|^2) \quad \text{for all iterations } t = 0, 1, 2, \dots \quad (25)$$

Thus, Corollary 1 guarantees that the optimization error decreases geometrically, with contraction factor κ , up to a tolerance that is of strictly lower order than the statistical precision $\|\hat{\theta} - \theta^*\|^2$. As will be clarified in several examples to follow, the condition $\Psi^2(\overline{\mathcal{M}})(\tau_u(\mathcal{L}_n) + \tau_\ell(\mathcal{L}_n)) = o(1)$ is satisfied for many statistical models, including sparse linear regression and low-rank matrix regression. This result is illustrated in Figure 2(b), where the solid circle represents the optimization tolerance, and the dotted circle represents the statistical precision. In the results to follow, we will quantify the term $o(\|\hat{\theta} - \theta^*\|^2)$ in a more precise manner for different statistical models.

Convergence rates for composite gradient: We now present our main result for the composite gradient iterates (4) that are suitable for the Lagrangian-based estimator (2). As before, our analysis yields a range of bounds indexed by subspace pairs $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$ that are \mathcal{R} -decomposable. For any subspace $\overline{\mathcal{M}}$ such that $64\tau_\ell(\mathcal{L}_n)\Psi^2(\overline{\mathcal{M}}) < \gamma_\ell$, we define *effective RSC coefficient* as

$$\overline{\gamma}_\ell := \gamma_\ell - 64\tau_\ell(\mathcal{L}_n)\Psi^2(\overline{\mathcal{M}}). \quad (26)$$

This coefficient accounts for the residual amount of strong convexity after accounting for the lower tolerance terms. In addition, we define the *compound contraction coefficient* as

$$\kappa(\mathcal{L}_n; \overline{\mathcal{M}}) := \left\{ 1 - \frac{\overline{\gamma}_\ell}{4\gamma_u} + \frac{64\Psi^2(\overline{\mathcal{M}})\tau_u(\mathcal{L}_n)}{\overline{\gamma}_\ell} \right\} \xi(\overline{\mathcal{M}}) \quad (27)$$

where $\xi(\overline{\mathcal{M}}) := (1 - \frac{64\tau_u(\mathcal{L}_n)\Psi^2(\overline{\mathcal{M}})}{\overline{\gamma}_\ell})^{-1}$, and $\Delta^* = \widehat{\theta}_{\lambda_n} - \theta^*$ is the statistical error vector⁵ for a specific choice of $\bar{\rho}$ and λ_n . As before, the coefficient κ measures the geometric rate of convergence for the algorithm. Finally, we define the *compound tolerance parameter*

$$\epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}}) := 8\xi(\overline{\mathcal{M}})\beta(\overline{\mathcal{M}})(6\Psi(\overline{\mathcal{M}})\|\Delta^*\| + 8\mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)))^2, \quad (28)$$

where $\beta(\overline{\mathcal{M}}) := 2\left(\frac{\overline{\gamma}_\ell}{4\gamma_u} + \frac{128\tau_u(\mathcal{L}_n)\Psi^2(\overline{\mathcal{M}})}{\overline{\gamma}_\ell}\right)\tau_\ell(\mathcal{L}_n) + 8\tau_u(\mathcal{L}_n) + 2\tau_\ell(\mathcal{L}_n)$. As with our previous result, the tolerance parameter determines the radius up to which geometric convergence can be attained.

Recall that the regularized problem (2) involves both a regularization weight λ_n , and a constraint radius $\bar{\rho}$. Our theory requires that the constraint radius is chosen such that $\bar{\rho} \geq \mathcal{R}(\theta^*)$, which ensures that θ^* is feasible. In addition, the regularization parameter should be chosen to satisfy the constraint

$$\lambda_n \geq 2\mathcal{R}^*(\nabla\mathcal{L}_n(\theta^*)), \quad (29)$$

where \mathcal{R}^* is the dual norm of the regularizer. This constraint is known to play an important role in proving bounds on the statistical error of regularized M -estimators (see the paper [28] and references therein for further details). Recalling the definition (2) of the overall objective function $\phi_n(\theta)$, the following result provides bounds on the *excess loss* $\phi_n(\theta^t) - \phi_n(\widehat{\theta}_{\lambda_n})$.

Theorem 2. *Consider the optimization problem (2) for a radius $\bar{\rho}$ such that θ^* is feasible, and a regularization parameter λ_n satisfying the bound (29), and suppose that the loss function \mathcal{L}_n satisfies the RSC/RSM condition with parameters $(\gamma_\ell, \tau_\ell(\mathcal{L}_n))$ and $(\gamma_u, \tau_u(\mathcal{L}_n))$ respectively. Let $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$ be any \mathcal{R} -decomposable pair such that*

$$\kappa \equiv \kappa(\mathcal{L}_n, \overline{\mathcal{M}}) \in [0, 1), \quad \text{and} \quad \frac{32\bar{\rho}}{1 - \kappa(\mathcal{L}_n; \overline{\mathcal{M}})}\xi(\overline{\mathcal{M}})\beta(\overline{\mathcal{M}}) \leq \lambda_n. \quad (30)$$

Then for any tolerance parameter $\delta^2 \geq \frac{\epsilon^2(\Delta^; \mathcal{M}, \overline{\mathcal{M}})}{(1-\kappa)}$, we have*

$$\phi_n(\theta^t) - \phi_n(\widehat{\theta}_{\lambda_n}) \leq \delta^2 \quad \text{for all } t \geq \frac{2\log \frac{\phi_n(\theta^0) - \phi_n(\widehat{\theta}_{\lambda_n})}{\delta^2}}{\log(1/\kappa)} + \log_2 \log_2 \left(\frac{\bar{\rho}\lambda_n}{\delta^2} \right) \left(1 + \frac{\log 2}{\log(1/\kappa)} \right). \quad (31)$$

⁵When the context is clear, we remind the reader that we drop the subscript λ_n on the parameter $\widehat{\theta}$.

Remarks: Note that the bound (31) guarantees the excess loss $\phi_n(\theta^t) - \phi_n(\hat{\theta})$ decays geometrically up to any squared error δ^2 larger than the compound tolerance (28). Moreover, the RSC condition also allows us to translate this bound on objective values to a bound on the optimization error $\theta^t - \hat{\theta}$. In particular, for any iterate θ^t such that $\phi_n(\theta^t) - \phi_n(\hat{\theta}) \leq \delta^2$, we are guaranteed that

$$\|\theta^t - \hat{\theta}_{\lambda_n}\|^2 \leq \frac{2\delta^2}{\gamma_\ell} + \frac{16\delta^2\tau_\ell(\mathcal{L}_n)}{\gamma_\ell\lambda_n^2} + \frac{4\tau_\ell(\mathcal{L}_n)(6\Psi(\bar{\mathcal{M}}) + 8\mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)))^2}{\gamma_\ell}. \quad (32)$$

In conjunction with Theorem 2, we see that it suffices to take a number of steps that is logarithmic in the inverse tolerance $(1/\delta)$, again showing a geometric rate of convergence.

Whereas Theorem 1 requires setting the radius so that the constraint is active, Theorem 2 has only a very mild constraint on the radius $\bar{\rho}$, namely that it be large enough such that $\bar{\rho} \geq \mathcal{R}(\theta^*)$. The reason for this much milder requirement is that the additive regularization with weight λ_n suffices to constrain the solution, whereas the extra side constraint is only needed to ensure good behavior of the optimization algorithm in the first few iterations. The regularization parameter λ_n must satisfy the so-called dual norm condition (29), which has appeared in past literature on statistical estimation, and is well-characterized for a broad range of statistical models (e.g., see the paper [28] and references therein).

Step-size setting: It seems that the updates (3) and (4) need to know the smoothness bound γ_u in order to set the step-size for gradient updates. However, we can use the same doubling trick as described in Algorithm (3.1) of Nesterov [32]. At each step, we check if the smoothness upper bound holds at the current iterate relative to the previous one. If the condition does not hold, we double our estimate of γ_u and resume. This guarantees a geometric convergence with a contraction factor worse at most by a factor of 2, compared to the knowledge of γ_u . We refer the reader to Nesterov [32] for details.

The following subsections are devoted to the development of some consequences of Theorems 1 and 2 and Corollary 1 for some specific statistical models, among them sparse linear regression with ℓ_1 -regularization, and matrix regression with nuclear norm regularization. In contrast to the entirely deterministic arguments that underlie the Theorems 1 and 2, these corollaries involve probabilistic arguments, more specifically in order to establish that the RSC and RSM properties hold with high probability.

3.2 Sparse vector regression

Recall from Section 2.4.1 the observation model for sparse linear regression. In a variety of applications, it is natural to assume that θ^* is sparse. For a parameter $q \in [0, 1]$ and radius $R_q > 0$, let us define the ℓ_q “ball”

$$\mathbb{B}_q(R_q) := \left\{ \theta \in \mathbb{R}^d \mid \sum_{j=1}^d |\beta_j|^q \leq R_q \right\}. \quad (33)$$

Note that $q = 0$ corresponds to the case of “hard sparsity”, for which any vector $\beta \in \mathbb{B}_0(R_0)$ is supported on a set of cardinality at most R_0 . For $q \in (0, 1]$, membership in the set $\mathbb{B}_q(R_q)$ enforces a decay rate on the ordered coefficients, thereby modelling approximate sparsity. In order to

estimate the unknown regression vector $\theta^* \in \mathbb{B}_q(R_q)$, we consider the least-squares Lasso estimator from Section 2.4.1, based on the quadratic loss function $\mathcal{L}(\theta; Z_1^n) := \frac{1}{2n} \|y - X\theta\|_2^2$, where $X \in \mathbb{R}^{n \times d}$ is the design matrix. In order to state a concrete result, we consider a random design matrix X , in which each row $x_i \in \mathbb{R}^d$ is drawn i.i.d. from a $N(0, \Sigma)$ distribution, where Σ is a positive definite covariance matrix. We refer to this as the *Σ -ensemble of random design matrices*, and use $\sigma_{\max}(\Sigma)$ and $\sigma_{\min}(\Sigma)$ to refer the maximum and minimum eigenvalues of Σ respectively, and $\zeta(\Sigma) := \max_{j=1,2,\dots,d} \Sigma_{jj}$ for the maximum variance. We also assume that the observation noise is zero-mean and sub-Gaussian with parameter ν^2 .

Guarantees for constrained Lasso: Our convergence rate on the optimization error $\theta^t - \hat{\theta}$ is stated in terms of the contraction coefficient

$$\kappa := \left\{ 1 - \frac{\sigma_{\min}(\Sigma)}{4\sigma_{\max}(\Sigma)} + \chi_n(\Sigma) \right\} \left\{ 1 - \chi_n(\Sigma) \right\}^{-1}, \quad (34)$$

where we have adopted the shorthand

$$\chi_n(\Sigma) := \begin{cases} \frac{c_0 \zeta(\Sigma)}{\sigma_{\max}(\Sigma)} R_q \left(\frac{\log d}{n} \right)^{1-q/2} & \text{for } q > 0 \\ \frac{c_0 \zeta(\Sigma)}{\sigma_{\max}(\Sigma)} s \left(\frac{\log d}{n} \right) & \text{for } q = 0 \end{cases}, \quad \text{for a numerical constant } c_0, \quad (35)$$

We assume that $\chi_n(\Sigma)$ is small enough to ensure that $\kappa \in (0, 1)$; in terms of the sample size, this amounts to a condition of the form $n = \Omega(R_q^{1/(1-q/2)} \log d)$. Such a scaling is sensible, since it is known from minimax theory on sparse linear regression [35] to be necessary for any method to be statistically consistent over the ℓ_q -ball.

With this set-up, we have the following consequence of Theorem 1:

Corollary 2 (Sparse vector recovery). *Under conditions of Theorem 1, suppose that we solve the constrained Lasso with $\rho \leq \|\theta^*\|_1$.*

(a) **Exact sparsity:** *If θ^* is supported on a subset of cardinality s , then with probability at least $1 - \exp(-c_1 \log d)$, the iterates (3) with $\gamma_u = 2\sigma_{\max}(\Sigma)$ satisfy*

$$\|\theta^t - \hat{\theta}\|_2^2 \leq \kappa^t \|\theta^0 - \hat{\theta}\|_2^2 + c_2 \chi_n(\Sigma) \|\hat{\theta} - \theta^*\|_2^2 \quad \text{for all } t = 0, 1, 2, \dots \quad (36)$$

(b) **Weak sparsity:** *Suppose that $\theta^* \in \mathbb{B}_q(R_q)$ for some $q \in (0, 1]$. Then with probability at least $1 - \exp(-c_1 \log d)$, the iterates (3) with $\gamma_u = 2\sigma_{\max}(\Sigma)$ satisfy*

$$\|\theta^t - \hat{\theta}\|_2^2 \leq \kappa^t \|\theta^0 - \hat{\theta}\|_2^2 + c_2 \chi_n(\Sigma) \left\{ R_q \left(\frac{\log d}{n} \right)^{1-q/2} + \|\hat{\theta} - \theta^*\|_2^2 \right\}. \quad (37)$$

We provide the proof of Corollary 2 in Section 5.4. Here we compare part (a), which deals with the special case of exactly sparse vectors, to some past work that has established convergence guarantees for optimization algorithms for sparse linear regression. Certain methods are known to converge at sublinear rates (e.g., [4]), more specifically at the rate $\mathcal{O}(1/t^2)$. The geometric rate of convergence guaranteed by Corollary 2 is exponentially faster. Other work on sparse regression has provided geometric rates of convergence that hold once the iterates are close to the optimum [8,

18], or geometric convergence up to the noise level ν^2 using various methods, including greedy methods [42] and thresholded gradient methods [17]. In contrast, Corollary 2 guarantees geometric convergence for all iterates up to a precision below that of statistical error. For these problems, the statistical error $\frac{\nu^2 s \log d}{n}$ is typically much smaller than the noise variance ν^2 , and decreases as the sample size is increased.

In addition, Corollary 2 also applies to the case of approximately sparse vectors, lying within the set $\mathbb{B}_q(R_q)$ for $q \in (0, 1]$. There are some important differences between the case of exact sparsity (Corollary 2(a)) and that of approximate sparsity (Corollary 2(b)). Part (a) guarantees geometric convergence to a tolerance depending only on the statistical error $\|\hat{\theta} - \theta^*\|_2$. In contrast, the second result also has the additional term $R_q(\frac{\log d}{n})^{1-q/2}$. This second term arises due to the statistical non-identifiability of linear regression over the ℓ_q -ball, and it is no larger than $\|\hat{\theta} - \theta^*\|_2^2$ with high probability. This assertion follows from known results [35] about minimax rates for linear regression over ℓ_q -balls; these unimprovable rates include a term of this order.

Guarantees for regularized Lasso: Using similar methods, we can also use Theorem 2 to obtain an analogous guarantee for the regularized Lasso estimator. Here focus only on the case of exact sparsity, although the result extends to approximate sparsity in a similar fashion. Letting $c_i, i = 0, 1, 2, 3, 4$ be universal positive constants, we define the modified curvature constant $\bar{\gamma}_\ell := \gamma_\ell - c_0 \frac{s \log d}{n} \zeta(\Sigma)$. Our results assume that $n = \Omega(s \log d)$, a condition known to be necessary for statistical consistency, so that $\bar{\gamma}_\ell > 0$. The contraction factor then takes the form

$$\kappa := \left\{ 1 - \frac{\sigma_{\min}(\Sigma)}{16\sigma_{\max}(\Sigma)} + c_1 \chi_n(\Sigma) \right\} \left\{ 1 - c_2 \chi_n(\Sigma) \right\}^{-1}, \quad \text{where } \chi_n(\Sigma) = \frac{\zeta(\Sigma)}{\bar{\gamma}_\ell} \frac{s \log d}{n}.$$

The tolerance factor in the optimization is given by

$$\epsilon_{\text{tol}}^2 := \frac{5 + c_2 \chi_n(\Sigma)}{1 - c_3 \chi_n(\Sigma)} \frac{\zeta(\Sigma) s \log d}{n} \|\theta^* - \hat{\theta}\|_2^2, \quad (38)$$

where $\theta^* \in \mathbb{R}^d$ is the unknown regression vector, and $\hat{\theta}$ is any optimal solution. With this notation, we have the following corollary.

Corollary 3 (Regularized Lasso). *Under conditions of Theorem 2, suppose that we solve the regularized Lasso with $\lambda_n = 6\sqrt{\frac{\nu \log d}{n}}$, and that θ^* is supported on a subset of cardinality at most s . Suppose that we have the condition*

$$64\bar{\rho} \frac{\log d}{n} \frac{5 + \frac{\bar{\gamma}_\ell}{4\gamma_u} + \frac{64s \log d/n}{\bar{\gamma}_\ell}}{\frac{\bar{\gamma}_\ell}{4\gamma_u} - \frac{128s \log d/n}{\bar{\gamma}_\ell}} \leq \lambda_n. \quad (39)$$

Then with probability at least $1 - \exp(-c_4 \log d)$, for any $\delta^2 \geq \epsilon_{\text{tol}}^2$, for any optimum $\hat{\theta}_{\lambda_n}$, we have

$$\|\theta^t - \hat{\theta}_{\lambda_n}\|_2^2 \leq \delta^2 \quad \text{for all iterations } t \geq (\log \frac{\phi_n(\theta^0) - \phi_n(\hat{\theta}_{\lambda_n})}{\delta^2}) / (\log \frac{1}{\kappa}).$$

As with Corollary 2(a), this result guarantees that $\mathcal{O}(\log(1/\epsilon_{\text{tol}}^2))$ iterations are sufficient to obtain an iterate θ^t that is within squared error $\mathcal{O}(\epsilon_{\text{tol}}^2)$ of any optimum $\hat{\theta}_{\lambda_n}$. The condition (39) is the specialization of Equation 30 to the sparse linear regression problem, and imposes an upper bound on admissible settings of $\bar{\rho}$ for our theory. Moreover, whenever $\frac{s \log d}{n} = o(1)$ —a condition that is required for statistical consistency of *any method*—the optimization tolerance ϵ_{tol}^2 is of lower order than the statistical error $\|\theta^* - \theta\|_2^2$.

3.3 Matrix regression with rank constraints

We now turn to estimation of matrices under various types of “soft” rank constraints. Recall the model of matrix regression from Section 2.4.2, and the M -estimator based on least-squares regularized with the nuclear norm (19). So as to reduce notational overhead, here we specialize to square matrices $\Theta^* \in \mathbb{R}^{d \times d}$, so that our observations are of the form

$$y_i = \langle\langle X_i, \Theta^* \rangle\rangle + w_i, \quad \text{for } i = 1, 2, \dots, n, \quad (40)$$

where $X_i \in \mathbb{R}^{d \times d}$ is a matrix of covariates, and $w_i \sim N(0, \nu^2)$ is Gaussian noise. As discussed in Section 2.4.2, the nuclear norm $\mathcal{R}(\Theta) = \|\Theta\|_1 = \sum_{j=1}^d \sigma_j(\Theta)$ is decomposable with respect to appropriately chosen matrix subspaces, and we exploit this fact heavily in our analysis.

We model the behavior of both exactly and approximately low-rank matrices by enforcing a sparsity condition on the vector $\sigma(\Theta) = [\sigma_1(\Theta) \ \sigma_2(\Theta) \ \dots \ \sigma_d(\Theta)]$ of singular values. In particular, for a parameter $q \in [0, 1]$, we define the ℓ_q -“ball” of matrices

$$\mathbb{B}_q(R_q) := \left\{ \Theta \in \mathbb{R}^{d \times d} \mid \sum_{j=1}^d |\sigma_j(\Theta)|^q \leq R_q \right\}. \quad (41)$$

Note that if $q = 0$, then $\mathbb{B}_0(R_0)$ consists of the set of all matrices with rank at most $r = R_0$. On the other hand, for $q \in (0, 1]$, the set $\mathbb{B}_q(R_q)$ contains matrices of all ranks, but enforces a relatively fast rate of decay on the singular values.

3.3.1 Bounds for matrix compressed sensing

We begin by considering the compressed sensing version of matrix regression, a model first introduced by Recht et al. [37], and later studied by other authors (e.g., [24, 29]). In this model, the observation matrices $X_i \in \mathbb{R}^{d \times d}$ are dense and drawn from some random ensemble. The simplest example is the standard Gaussian ensemble, in which each entry of X_i is drawn i.i.d. as standard normal $N(0, 1)$. Note that X_i is a dense matrix in general; this is an important contrast with the matrix completion setting to follow shortly.

Here we consider a more general ensemble of random matrices X_i , in which each matrix $X_i \in \mathbb{R}^{d \times d}$ is drawn i.i.d. from a zero-mean normal distribution in \mathbb{R}^{d^2} with covariance matrix $\Sigma \in \mathbb{R}^{d^2 \times d^2}$. The setting $\Sigma = I_{d^2 \times d^2}$ recovers the standard Gaussian ensemble studied in past work. As usual, we let $\sigma_{\max}(\Sigma)$ and $\sigma_{\min}(\Sigma)$ define the maximum and minimum eigenvalues of Σ , and we define $\zeta_{\text{mat}}(\Sigma) = \sup_{\|u\|_2=1} \sup_{\|v\|_2=1} \text{var}(\langle\langle X, uv^T \rangle\rangle)$, corresponding to the maximal variance of X when projected onto rank one matrices. For the identity ensemble, we have $\zeta_{\text{mat}}(I) = 1$.

We now state a result on the convergence of the updates (20) when applied to a statistical problem involving a matrix $\Theta^* \in \mathbb{B}_q(R_q)$. The convergence rate depends on the contraction coefficient

$$\kappa := \left\{ 1 - \frac{\sigma_{\min}(\Sigma)}{4\sigma_{\max}(\Sigma)} + \chi_n(\Sigma) \right\} \left\{ 1 - \chi_n(\Sigma) \right\}^{-1},$$

where $\chi_n(\Sigma) := \frac{c_1 \zeta_{\text{mat}}(\Sigma)}{\sigma_{\max}(\Sigma)} R_q \left(\frac{d}{n} \right)^{1-q/2}$ for some universal constant c_1 . In the case $q = 0$, corresponding to matrices with rank at most r , note that we have $R_0 = r$. With this notation, we have the following convergence guarantee:

Corollary 4 (Low-rank matrix recovery). *Under conditions of Theorem 1, consider the semidefinite program (19) with $\rho \leq \|\Theta^*\|_1$, and suppose that we apply the projected gradient updates (20) with $\gamma_u = 2\sigma_{\max}(\Sigma)$.*

- (a) Exactly low-rank: *In the case $q = 0$, if Θ^* has rank $r < d$, then with probability at least $1 - \exp(-c_0 d)$, the iterates (20) satisfy the bound*

$$\|\Theta^t - \widehat{\Theta}\|_F^2 \leq \kappa^t \|\Theta^0 - \widehat{\Theta}\|_F^2 + c_2 \chi_n(\Sigma) \|\widehat{\Theta} - \Theta^*\|_F^2 \quad \text{for all } t = 0, 1, 2, \dots \quad (42)$$

- (b) Approximately low-rank: *If $\Theta^* \in \mathbb{B}_q(R_q)$ for some $q \in (0, 1]$, then with probability at least $1 - \exp(-c_0 d)$, the iterates (20) satisfy*

$$\|\Theta^t - \widehat{\Theta}\|_F^2 \leq \kappa^t \|\Theta^0 - \widehat{\Theta}\|_F^2 + c_2 \chi_n(\Sigma) \left\{ R_q \left(\frac{d}{n} \right)^{1-q/2} + \|\widehat{\Theta} - \Theta^*\|_F^2 \right\}, \quad (43)$$

Although quantitative aspects of the rates are different, Corollary 4 is analogous to Corollary 2. For the case of exactly low rank matrices (part (a)), geometric convergence is guaranteed up to a tolerance involving the statistical error $\|\widehat{\Theta} - \Theta^*\|_F^2$. For the case of approximately low rank matrices (part (b)), the tolerance term involves an additional factor of $R_q \left(\frac{d}{n} \right)^{1-q/2}$. Again, from known results on minimax rates for matrix estimation [38], this term is known to be of comparable or lower order than the quantity $\|\widehat{\Theta} - \Theta^*\|_F^2$. As before, it is also possible to derive an analogous corollary of Theorem 2 for estimating low-rank matrices; in the interests of space, we leave such a development to the reader.

3.3.2 Bounds for matrix completion

In this model, observation y_i is a noisy version of a randomly selected entry $\Theta_{a(i),b(i)}^*$ of the unknown matrix Θ^* . Applications of this matrix completion problem include collaborative filtering [40], where the rows of the matrix Θ^* correspond to users, and the columns correspond to items (e.g., movies in the Netflix database), and the entry Θ_{ab}^* corresponds to user's a rating of item b . Given observations of only a subset of the entries of Θ^* , the goal is to fill in, or complete the matrix, thereby making recommendations of movies that a given user has not yet seen.

Matrix completion can be viewed as a particular case of the matrix regression model (18), in particular by setting $X_i = E_{a(i)b(i)}$, corresponding to the matrix with a single one in position $(a(i), b(i))$, and zeroes in all other positions. Note that these observation matrices are extremely sparse, in contrast to the compressed sensing model. Nuclear-norm based estimators for matrix completion are known to have good statistical properties (e.g., [11, 36, 40, 30]). Here we consider the M -estimator

$$\widehat{\Theta} \in \arg \min_{\Theta \in \Omega} \frac{1}{2n} \sum_{i=1}^n (y_i - \Theta_{a(i)b(i)})^2 \quad \text{such that } \|\Theta\|_1 \leq \rho, \quad (44)$$

where $\Omega = \{\Theta \in \mathbb{R}^{d \times d} \mid \|\Theta\|_\infty \leq \frac{\alpha}{d}\}$ is the set of matrices with bounded elementwise ℓ_∞ norm. This constraint eliminates matrices that are overly “spiky” (i.e., concentrate too much of their mass in a single position); as discussed in the paper [30], such spikiness control is necessary in order to bound the non-identifiable component of the matrix completion model.

Corollary 5 (Matrix completion). *Under the conditions of Theorem 1, suppose that $\Theta^* \in \mathbb{B}_q(R_q)$, and that we solve the program (44) with $\rho \leq \|\Theta^*\|_1$. As long as $n > c_0 R_q^{1/(1-q/2)} d \log d$ for a sufficiently large constant c_0 , then with probability at least $1 - \exp(-c_1 d \log d)$, there is a contraction coefficient $\kappa_t \in (0, 1)$ that decreases with t such that for all iterations $t = 0, 1, 2, \dots$,*

$$\|\Theta^{t+1} - \widehat{\Theta}\|_F^2 \leq \kappa_t^t \|\Theta^0 - \widehat{\Theta}\|_F^2 + c_2 \left\{ R_q \left(\frac{\alpha^2 d \log d}{n} \right)^{1-q/2} + \|\widehat{\Theta} - \Theta^*\|_F^2 \right\}. \quad (45)$$

In some cases, the bound on $\|\Theta\|_\infty$ in the algorithm (44) might be unknown, or undesirable. While this constraint is necessary in general [30], it can be avoided if more information such as the sampling distribution (that is, the distribution of X_i) is known and used to construct the estimator. In this case, Koltchinskii et al. [22] show error bounds on a nuclear norm penalized estimator without requiring ℓ_∞ bound on $\widehat{\Theta}$.

Again a similar corollary of Theorem 2 can be derived by combining the proof of Corollary 5 with that of Theorem 2. An interesting aspect of this problem is that the condition 30(b) takes the form $\lambda_n > \frac{c\alpha\sqrt{d \log d/n}}{1-\kappa}$, where α is a bound on $\|\Theta\|_\infty$. This condition is independent of $\bar{\rho}$, and hence, given a sample size as stated in the corollary, the algorithm always converges geometrically for any radius $\bar{\rho} \geq \|\Theta^*\|_1$.

3.4 Matrix decomposition problems

In recent years, various researchers have studied methods for solving the problem of matrix decomposition (e.g., [12, 10, 44, 1, 19]). The basic problem has the following form: given a pair of unknown matrices Θ^* and Γ^* , both lying in $\mathbb{R}^{d_1 \times d_2}$, suppose that we observe a third matrix specified by the model $Y = \Theta^* + \Gamma^* + W$, where $W \in \mathbb{R}^{d_1 \times d_2}$ represents observation noise. Typically the matrix Θ^* is assumed to be low-rank, and some low-dimensional structural constraint is assumed on the matrix Γ^* . For example, the papers [12, 10, 19] consider the setting in which Γ^* is sparse, while Xu et al. [44] consider a column-sparse model, in which only a few of the columns of Γ^* have non-zero entries. In order to illustrate the application of our general result to this setting, here we consider the low-rank plus column-sparse framework [44]. (We note that since the ℓ_1 -norm is decomposable, similar results can easily be derived for the low-rank plus entrywise-sparse setting as well.)

Since Θ^* is assumed to be low-rank, as before we use the nuclear norm $\|\Theta\|_1$ as a regularizer (see Section 2.4.2). We assume that the unknown matrix $\Gamma^* \in \mathbb{R}^{d_1 \times d_2}$ is column-sparse, say with at most $s < d_2$ non-zero columns. A suitable convex regularizer for this matrix structure is based on the *columnwise* $(1, 2)$ -norm, given by

$$\|\Gamma\|_{1,2} := \sum_{j=1}^{d_2} \|\Gamma_j\|_2, \quad (46)$$

where $\Gamma_j \in \mathbb{R}^{d_1}$ denotes the j^{th} column of Γ . Note also that the dual norm is given by the *elementwise* $(\infty, 2)$ -norm $\|\Gamma\|_{\infty,2} = \max_{j=1,\dots,d_2} \|\Gamma_j\|_2$, corresponding to the maximum ℓ_2 -norm over columns.

In order to estimate the unknown pair (Θ^*, Γ^*) , we consider the M -estimator

$$(\widehat{\Theta}, \widehat{\Gamma}) := \arg \min_{\Theta, \Gamma} \|Y - \Theta - \Gamma\|_F^2 \quad \text{such that} \quad \|\Theta\|_1 \leq \rho_\Theta, \quad \|\Gamma\|_{1,2} \leq \rho_\Gamma \text{ and } \|\Theta\|_{\infty,2} \leq \frac{\alpha}{\sqrt{d_2}} \quad (47)$$

The first two constraints restrict Θ and Γ to a nuclear norm ball of radius ρ_Θ and a $(1, 2)$ -norm ball of radius ρ_Γ , respectively. The final constraint controls the “spikiness” of the low-rank component Θ , as measured in the $(\infty, 2)$ -norm, corresponding to the maximum ℓ_2 -norm over the columns. As with the elementwise ℓ_∞ -bound for matrix completion, this additional constraint is required in order to limit the non-identifiability in matrix decomposition. (See the paper [1] for more discussion of non-identifiability issues in matrix decomposition.)

With this set-up, consider the projected gradient algorithm when applied to the matrix decomposition problem: it generates a sequence of matrix pairs (Θ^t, Γ^t) for $t = 0, 1, 2, \dots$, and the optimization error is characterized in terms of the matrices $\widehat{\Delta}_\Theta^t := \Theta^t - \widehat{\Theta}$ and $\widehat{\Delta}_\Gamma^t := \Gamma^t - \widehat{\Gamma}$. Finally, we measure the optimization error at time t in terms of the squared Frobenius error $e^2(\widehat{\Delta}_\Theta^t, \widehat{\Delta}_\Gamma^t) := \|\widehat{\Delta}_\Theta^t\|_F^2 + \|\widehat{\Delta}_\Gamma^t\|_F^2$, summed across both the low-rank and column-sparse components.

Corollary 6 (Matrix decomposition). *Under the conditions of Theorem 1, suppose that $\|\Theta^*\|_{\infty, 2} \leq \frac{\alpha}{\sqrt{d_2}}$ and Γ^* has at most s non-zero columns. If we solve the convex program (47) with $\rho_\Theta \leq \|\Theta^*\|_1$ and $\rho_\Gamma \leq \|\Gamma^*\|_{1,2}$, then for all iterations $t = 0, 1, 2, \dots$,*

$$e^2(\widehat{\Delta}_\Theta^t, \widehat{\Delta}_\Gamma^t) \leq \left(\frac{3}{4}\right)^t e^2(\widehat{\Delta}_\Theta^0, \widehat{\Delta}_\Gamma^0) + c \left(\|\widehat{\Gamma} - \Gamma^*\|_F^2 + \alpha^2 \frac{s}{d_2} \right).$$

This corollary has some unusual aspects, relative to the previous corollaries. First of all, in contrast to the previous results, the guarantee is a deterministic one (as opposed to holding with high probability). More specifically, the RSC/RSM conditions hold deterministic sense, which should be contrasted with the high probability statements given in Corollaries 2-5. Consequently, the effective conditioning of the problem does not depend on sample size and we are guaranteed geometric convergence at a fixed rate, independent of sample size. The additional tolerance term is completely independent of the rank of Θ^* and only depends on the column-sparsity of Γ^* .

4 Simulation results

In this section, we provide some experimental results that confirm the accuracy of our theoretical results, in particular showing excellent agreement with the linear rates predicted by our theory. In addition, the rates of convergence slow down for smaller sample sizes, which lead to problems with relatively poor conditioning. In all the simulations reported below, we plot the log error $\|\theta^t - \widehat{\theta}\|$ between the iterate θ^t at time t versus the final solution $\widehat{\theta}$. Each curve provides the results averaged over five random trials, according to the ensembles which we now describe.

4.1 Sparse regression

We begin by considering the linear regression model $y = X\theta^* + w$ where θ^* is the unknown regression vector belonging to the set $\mathbb{B}_q(R_q)$, and i.i.d. observation noise $w_i \sim N(0, 0.25)$. We consider a family of ensembles for the random design matrix $X \in \mathbb{R}^{n \times d}$. In particular, we construct X by generating each row $x_i \in \mathbb{R}^d$ independently according to following procedure. Let z_1, \dots, z_n be an i.i.d. sequence of $N(0, 1)$ variables, and fix some correlation parameter $\omega \in [0, 1)$. We first initialize by setting $x_{i,1} = z_1 / \sqrt{1 - \omega^2}$, and then generate the remaining entries by applying the recursive update $x_{i,t+1} = \omega x_{i,t} + z_t$ for $t = 1, 2, \dots, d - 1$, so that $x_i \in \mathbb{R}^d$ is a zero-mean Gaussian

random vector. It can be verified that all the eigenvalues of $\Sigma = \text{cov}(x_i)$ lie within the interval $[\frac{1}{(1+\omega)^2}, \frac{2}{(1-\omega)^2(1+\omega)}]$, so that Σ has a finite condition number for all $\omega \in [0, 1)$. At one extreme, for $\omega = 0$, the matrix Σ is the identity, and so has condition number equal to 1. As $\omega \rightarrow 1$, the matrix Σ becomes progressively more ill-conditioned, with a condition number that is very large for ω close to one. As a consequence, although incoherence conditions like the restricted isometry property can be satisfied when $\omega = 0$, they will fail to be satisfied (w.h.p.) once ω is large enough.

For this random ensemble of problems, we have investigated convergence rates for a wide range of dimensions d and radii R_q . Since the results are relatively uniform across the choice of these parameters, here we report results for dimension $d = 20,000$, and radius $R_q = \lceil (\log d)^2 \rceil$. In the case $q = 0$, the radius $R_0 = s$ corresponds to the sparsity level. The per iteration cost in this case is $\mathcal{O}(nd)$. In order to reveal dependence of convergence rates on sample size, we study a range of the form $n = \lceil \alpha s \log d \rceil$, where the *order parameter* $\alpha > 0$ is varied.

Our first experiment is based on taking the correlation parameter $\omega = 0$, and the ℓ_q -ball parameter $q = 0$, corresponding to exact sparsity. We then measure convergence rates for sample sizes specified by $\alpha \in \{1, 1.25, 5, 25\}$. As shown by the results plotted in panel (a) of Figure 3, projected gradient descent fails to converge for $\alpha = 1$ or $\alpha = 1.25$; in both these cases, the sample size n is too small for the RSC and RSM conditions to hold, so that a constant step size leads to oscillatory behavior in the algorithm. In contrast, once the order parameter α becomes large enough to ensure that the RSC/RSM conditions hold (w.h.p.), we observe a geometric convergence of the error $\|\theta^t - \hat{\theta}\|_2$. Moreover the convergence rate is faster for $\alpha = 25$ compared to $\alpha = 5$, since the RSC/RSM constants are better with larger sample size. Such behavior is in agreement with the conclusions of Corollary 2, which predicts that the convergence rate should improve as the number of samples n is increased.




Figure 3. Plot of the log of the optimization error $\log(\|\theta^t - \hat{\theta}\|_2)$ in the sparse linear regression problem, rescaled so the plots start at 0. In this problem, $d = 20000$, $s = \lceil \log d \rceil$, $n = \alpha s \log d$. Plot (a) shows convergence for the exact sparse case with $q = 0$ and $\Sigma = I$ (i.e. $\omega = 0$). In panel (b), we observe how convergence rates change as the correlation parameter ω is varied for $q = 0$ and $\alpha = 25$. Plot (c) shows the convergence rates when $\omega = 0$, $\alpha = 25$ and q is varied.

On the other hand, Corollary 2 also predicts that convergence rates should be slower when the condition number of Σ is worse. In order to test this prediction, we again studied an exactly sparse problem ($q = 0$), this time with the fixed sample size $n = \lceil 25s \log d \rceil$, and we varied the correlation parameter $\omega \in \{0, 0.5, 0.8\}$. As shown in panel (b) of Figure 3, the convergence rates slow down

as the correlation parameter is increased and for the case of extremely high correlation of $\omega = 0.8$, the optimization error curve is almost flat—the method makes very slow progress in this case.

A third prediction of Corollary 2 is that the convergence of projected gradient descent should become slower as the sparsity parameter q is varied between exact sparsity ($q = 0$), and the least sparse case ($q = 1$). (In particular, note for $n > \log d$, the quantity χ_n from equation (35) is monotonically increasing with q .) Panel (c) of Figure 3 shows convergence rates for the fixed sample size $n = 25s \log d$ and correlation parameter $\omega = 0$, and with the sparsity parameter $q \in \{0, 0.5, 1.0\}$. As expected, the convergence rate slows down as q increases from 0 to 1. Corollary 2 further captures how the contraction factor changes as the problem parameters (s, d, n) are varied. In particular, it predicts that as we change the triplet simultaneously, while holding the ratio $\alpha = s \log d/n$ constant, the convergence rate should stay the same. We recall that this phenomenon was indeed demonstrated in Figure 1 in Section 1.

4.2 Low-rank matrix estimation

We also performed experiments with two different versions of low-rank matrix regression. Our simulations applied to instances of the observation model $y_i = \langle\langle X_i, \Theta^* \rangle\rangle + w_i$, for $i = 1, 2, \dots, n$, where $\Theta^* \in \mathbb{R}^{200 \times 200}$ is a fixed unknown matrix, $X_i \in \mathbb{R}^{200 \times 200}$ is a matrix of covariates, and $w_i \sim N(0, 0.25)$ is observation noise. In analogy to the sparse vector problem, we performed simulations with the matrix Θ^* belonging to the set $\mathbb{B}_q(R_q)$ of approximately low-rank matrices, as previously defined in equation (41) for $q \in [0, 1]$. The case $q = 0$ corresponds to the set of matrices with rank at most $r = R_0$, whereas the case $q = 1$ corresponds to the ball of matrices with nuclear norm at most R_1 .




Figure 4. (a) Plot of $\log(\|\Theta^t - \hat{\Theta}\|_F)$ versus number of iterations in matrix compressed sensing for a matrix size $d = 200$ with rank $R_0 = 5$, and sample sizes $n = \alpha R_0 d$. For $\alpha \in \{1, 1.25\}$, the algorithm oscillates, whereas geometric convergence is obtained for $\alpha \in \{5, 25\}$, consistent with the theoretical prediction. (b) Plot of $\log(\|\Theta^t - \hat{\Theta}\|_F)$ versus number of iterations in matrix completion with $d = 200$, $R_0 = 5$, and $n = \alpha R_0 d \log(d)$ with $\alpha \in \{1, 2, 5, 25\}$. For $\alpha \in \{2, 5, 25\}$ the algorithm enjoys geometric convergence.

In our first set of matrix experiments, we considered the matrix version of compressed sensing [36], in which each matrix $X_i \in \mathbb{R}^{200 \times 200}$ is randomly formed with i.i.d. $N(0, 1)$ entries, as described in Section 3.3.1. In the case $q = 0$, we formed a matrix $\Theta^* \in \mathbb{R}^{200 \times 200}$ with rank $R_0 = 5$, and performed simulations over the sample sizes $n = \alpha R_0 d$, with the parameter $\alpha \in \{1, 1.25, 5, 25\}$. The per iteration cost in this case is $\mathcal{O}(nd^2)$. As seen in panel (a) of Figure 4, the projected gradient descent method exhibits behavior that is qualitatively similar to that for the sparse linear regression problem. More specifically, it fails to converge when the sample size (as reflected by the order parameter α) is too small, and converges geometrically with a progressively faster rate as α is increased. We have also observed similar types of scaling as the matrix sparsity parameter is increased from $q = 0$ to $q = 1$.

In our second set of matrix experiments, we studied the behavior of projected gradient descent for the problem of matrix completion, as described in Section 3.3.2. For this problem, we again studied matrices of dimension $d = 200$ and rank $R_0 = 5$, and we varied the sample size as $n = \alpha R_0 d \log d$ for $\alpha \in \{1, 2, 5, 25\}$. As shown in panel (b) of Figure 4, projected gradient descent for matrix completion also enjoys geometric convergence for α large enough.

5 Proofs

In this section, we provide the proofs of our results. Recall that we use $\widehat{\Delta}^t := \theta^t - \widehat{\theta}$ to denote the optimization error, and $\Delta^* = \widehat{\theta} - \theta^*$ to denote the statistical error. For future reference, we point out a slight weakening of restricted strong convexity (RSC), useful for obtaining parts of our results. As the proofs to follow reveal, it is only necessary to enforce an RSC condition of the form

$$\mathcal{T}_{\mathcal{L}}(\theta^t; \widehat{\theta}) \geq \frac{\gamma_{\ell}}{2} \|\theta^t - \widehat{\theta}\|^2 - \tau_{\ell}(\mathcal{L}_n) \mathcal{R}^2(\theta^t - \widehat{\theta}) - \delta^2, \quad (48)$$

which is milder than the original RSC condition (8), in that it applies only to differences of the form $\theta^t - \widehat{\theta}$, and allows for additional slack δ . We make use of this refined notion in the proofs of various results to follow.

With this relaxed RSC condition and the same RSM condition as before, our proof shows that

$$\|\theta^{t+1} - \widehat{\theta}\|^2 \leq \kappa^t \|\theta^0 - \widehat{\theta}\|^2 + \frac{\epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}}) + 2\delta^2/\gamma_u}{1 - \kappa} \quad \text{for all iterations } t = 0, 1, 2, \dots \quad (49)$$

Note that this result reduces to the previous statement when $\delta = 0$. This extension of Theorem 1 is used in the proofs of Corollaries 5 and 6.

We will assume without loss of generality that all the iterates lie in the subset Ω' of Ω . This can be ensured by augmenting the loss with the indicator of Ω' or equivalently performing projections on the set $\Omega' \cap \mathbb{B}_{\mathcal{R}}(\rho)$ as mentioned earlier.

5.1 Proof of Theorem 1

Recall that Theorem 1 concerns the constrained problem (1). The proof is based on two technical lemmas. The first lemma guarantees that at each iteration $t = 0, 1, 2, \dots$, the optimization error $\widehat{\Delta}^t = \theta^t - \widehat{\theta}$ belongs to an interesting constraint set defined by the regularizer.

Lemma 1. Let $\hat{\theta}$ be any optimum of the constrained problem (1) for which $\mathcal{R}(\hat{\theta}) = \rho$. Then for any iteration $t = 1, 2, \dots$ and for any \mathcal{R} -decomposable subspace pair $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$, the optimization error $\hat{\Delta}^t := \theta^t - \hat{\theta}$ belongs to the set

$$\mathbb{S}(\mathcal{M}; \overline{\mathcal{M}}; \theta^*) := \left\{ \Delta \in \Omega \mid \mathcal{R}(\Delta) \leq 2\Psi(\overline{\mathcal{M}}) \|\Delta\| + 2\mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)) + 2\mathcal{R}(\Delta^*) + \Psi(\overline{\mathcal{M}}) \|\Delta^*\| \right\}. \quad (50)$$

The proof of this lemma, provided in Appendix A.1, exploits the decomposability of the regularizer in an essential way.

The structure of the set (50) takes a simpler form in the special case when \mathcal{M} is chosen to contain θ^* and $\overline{\mathcal{M}} = \mathcal{M}$. In this case, we have $\mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)) = 0$, and hence the optimization error $\hat{\Delta}^t$ satisfies the inequality

$$\mathcal{R}(\hat{\Delta}^t) \leq 2\Psi(\mathcal{M}) \{ \|\hat{\Delta}^t\| + \|\Delta^*\| \} + 2\mathcal{R}(\Delta^*). \quad (51)$$

An inequality of this type, when combined with the definitions of RSC/RSM, allows us to establish the curvature conditions required to prove globally geometric rates of convergence.

We now state a second lemma under the more general RSC condition (48):

Lemma 2. Under the RSC condition (48) and RSM condition (10), for all $t = 0, 1, 2, \dots$, we have

$$\begin{aligned} & \gamma_u \langle \theta^t - \theta^{t+1}, \theta^t - \hat{\theta} \rangle \\ & \geq \left\{ \frac{\gamma_u}{2} \|\theta^t - \theta^{t+1}\|^2 - \tau_u(\mathcal{L}_n) \mathcal{R}^2(\theta^{t+1} - \theta^t) \right\} + \left\{ \frac{\gamma_\ell}{2} \|\theta^t - \hat{\theta}\|^2 - \tau_\ell(\mathcal{L}_n) \mathcal{R}^2(\theta^t - \hat{\theta}) - \delta^2 \right\}. \end{aligned} \quad (52)$$

The proof of this lemma, provided in Appendix A.2, follows along the lines of the intermediate result within Theorem 2.2.8 of Nesterov [31], but with some care required to handle the additional terms that arise in our weakened forms of strong convexity and smoothness.

Using these auxiliary results, let us now complete the the proof of Theorem 1. We first note the elementary relation

$$\|\theta^{t+1} - \hat{\theta}\|^2 = \|\theta^t - \hat{\theta} - \theta^t + \theta^{t+1}\|^2 = \|\theta^t - \hat{\theta}\|^2 + \|\theta^t - \theta^{t+1}\|^2 - 2\langle \theta^t - \hat{\theta}, \theta^t - \theta^{t+1} \rangle. \quad (53)$$

We now use Lemma 2 and the more general form of RSC (48) to control the cross-term, thereby obtaining the upper bound

$$\begin{aligned} \|\theta^{t+1} - \hat{\theta}\|^2 & \leq \|\theta^t - \hat{\theta}\|^2 - \frac{\gamma_\ell}{\gamma_u} \|\theta^t - \hat{\theta}\|^2 + \frac{2\tau_u(\mathcal{L}_n)}{\gamma_u} \mathcal{R}^2(\theta^{t+1} - \theta^t) + \frac{2\tau_\ell(\mathcal{L}_n)}{\gamma_u} \mathcal{R}^2(\theta^t - \hat{\theta}) + \frac{2\delta^2}{\gamma_u} \\ & = \left(1 - \frac{\gamma_\ell}{\gamma_u}\right) \|\theta^t - \hat{\theta}\|^2 + \frac{2\tau_u(\mathcal{L}_n)}{\gamma_u} \mathcal{R}^2(\theta^{t+1} - \theta^t) + \frac{2\tau_\ell(\mathcal{L}_n)}{\gamma_u} \mathcal{R}^2(\theta^t - \hat{\theta}) + \frac{2\delta^2}{\gamma_u}. \end{aligned}$$

We now observe that by triangle inequality and the Cauchy-Schwarz inequality,

$$\mathcal{R}^2(\theta^{t+1} - \theta^t) \leq (\mathcal{R}(\theta^{t+1} - \hat{\theta}) + \mathcal{R}(\hat{\theta} - \theta^t))^2 \leq 2\mathcal{R}^2(\theta^{t+1} - \hat{\theta}) + 2\mathcal{R}^2(\theta^t - \hat{\theta}).$$

Recall the definition of the optimization error $\widehat{\Delta}^t := \theta^t - \widehat{\theta}$, we have the upper bound

$$\|\widehat{\Delta}^{t+1}\|^2 \leq \left(1 - \frac{\gamma_\ell}{\gamma_u}\right)\|\widehat{\Delta}^t\|^2 + \frac{4\tau_u(\mathcal{L}_n)}{\gamma_u}\mathcal{R}^2(\widehat{\Delta}^{t+1}) + \frac{4\tau_u(\mathcal{L}_n) + 2\tau_\ell(\mathcal{L}_n)}{\gamma_u}\mathcal{R}^2(\widehat{\Delta}^t) + \frac{2\delta^2}{\gamma_u}. \quad (54)$$

We now apply Lemma 1 to control the terms involving \mathcal{R}^2 . In terms of squared quantities, the inequality (50) implies that

$$\mathcal{R}^2(\widehat{\Delta}^t) \leq 4\Psi^2(\overline{\mathcal{M}}^\perp)\|\widehat{\Delta}^t\|^2 + 2\nu^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}}) \quad \text{for all } t = 0, 1, 2, \dots,$$

where we recall that $\Psi^2(\overline{\mathcal{M}}^\perp)$ is the subspace compatibility (12) and $\nu^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}})$ accumulates all the residual terms. Applying this bound twice—once for t and once for $t+1$ —and substituting into equation (54) yields that $\left\{1 - \frac{16\Psi^2(\overline{\mathcal{M}}^\perp)\tau_u(\mathcal{L}_n)}{\gamma_u}\right\}\|\Delta^{t+1}\|^2$ is upper bounded by

$$\left\{1 - \frac{\gamma_\ell}{\gamma_u} + \frac{16\Psi^2(\overline{\mathcal{M}}^\perp)(\tau_u(\mathcal{L}_n) + \tau_\ell(\mathcal{L}_n))}{\gamma_u}\right\}\|\Delta^t\|^2 + \frac{16(\tau_u(\mathcal{L}_n) + \tau_\ell(\mathcal{L}_n))\nu^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}})}{\gamma_u} + \frac{2\delta^2}{\gamma_u}.$$

Under the assumptions of Theorem 1, we are guaranteed that $\frac{16\Psi^2(\overline{\mathcal{M}}^\perp)\tau_u(\mathcal{L}_n)}{\gamma_u} < 1/2$, and so we can re-arrange this inequality into the form

$$\|\Delta^{t+1}\|^2 \leq \kappa\|\Delta^t\|^2 + \epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}}) + \frac{2\delta^2}{\gamma_u} \quad (55)$$

where κ and $\epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}})$ were previously defined in equations (22) and (23) respectively. Iterating this recursion yields

$$\|\Delta^{t+1}\|^2 \leq \kappa^t\|\Delta^0\|^2 + \left(\epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}}) + \frac{2\delta^2}{\gamma_u}\right)\left(\sum_{j=0}^t \kappa^j\right).$$

The assumptions of Theorem 1 guarantee that $\kappa \in (0, 1)$, so that summing the geometric series yields the claim (24).

5.2 Proof of Theorem 2

The Lagrangian version of the optimization program is based on solving the convex program (2), with the objective function $\phi(\theta) = \mathcal{L}_n(\theta) + \lambda_n\mathcal{R}(\theta)$. Our proof is based on analyzing the error $\phi(\theta^t) - \phi(\widehat{\theta})$ as measured in terms of this objective function. It requires two technical lemmas, both of which are stated in terms of a given tolerance $\bar{\eta} > 0$, and an integer $T > 0$ such that

$$\phi(\theta^t) - \phi(\widehat{\theta}) \leq \bar{\eta} \quad \text{for all } t \geq T. \quad (56)$$

Our first technical lemma is analogous to Lemma 1, and restricts the optimization error $\widehat{\Delta}^t = \theta^t - \widehat{\theta}$ to a cone-like set.

Lemma 3 (Iterated Cone Bound (ICB)). *Let $\widehat{\theta}$ be any optimum of the regularized M -estimator (2). Under condition (56) with parameters $(T, \bar{\eta})$, for any iteration $t \geq T$ and for any \mathcal{R} -decomposable subspace pair $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$, the optimization error $\widehat{\Delta}^t := \theta^t - \widehat{\theta}$ satisfies*

$$\mathcal{R}(\widehat{\Delta}^t) \leq 4\Psi(\overline{\mathcal{M}})\|\widehat{\Delta}^t\| + 8\Psi(\overline{\mathcal{M}})\|\Delta^*\| + 8\mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)) + 2\min\left(\frac{\bar{\eta}}{\lambda_n}, \bar{\rho}\right) \quad (57)$$

Our next lemma guarantees sufficient decrease of the objective value difference $\phi(\theta^t) - \phi(\hat{\theta})$. Lemma 3 plays a crucial role in its proof. Recall the definition (27) of the compound contraction coefficient $\kappa(\mathcal{L}_n; \bar{\mathcal{M}})$, defined in terms of the related quantities $\xi(\bar{\mathcal{M}})$ and $\beta(\bar{\mathcal{M}})$. Throughout the proof, we drop the arguments of κ , ξ and β so as to ease notation.

Lemma 4. *Under the RSC (48) and RSM conditions (10), as well as assumption (56) with parameters $(\bar{\eta}, T)$, for all $t \geq T$, we have*

$$\phi(\theta^t) - \phi(\hat{\theta}) \leq \kappa^{t-T}(\phi(\theta^T) - \phi(\hat{\theta})) + \frac{2}{1-\kappa}\xi(\mathcal{M})\beta(\mathcal{M})(\varepsilon^2 + \bar{\epsilon}_{\text{stat}}^2),$$

where $\varepsilon := 2 \min(\bar{\eta}/\lambda_n, \bar{\rho})$ and $\bar{\epsilon}_{\text{stat}} := 8\Psi(\bar{\mathcal{M}})\|\Delta^*\| + 8\mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*))$.

We are now in a position to prove our main theorem, in particular via a recursive application of Lemma 4. At a high level, we divide the iterations $t = 0, 1, 2, \dots$ into a series of disjoint epochs $[T_k, T_{k+1})$ with $0 = T_0 \leq T_1 \leq T_2 \leq \dots$. Moreover, we define an associated sequence of tolerances $\bar{\eta}_0 > \bar{\eta}_1 > \dots$ such that at the end of epoch $[T_{k-1}, T_k)$, the optimization error has been reduced to $\bar{\eta}_k$. Our analysis guarantees that $\phi(\theta^t) - \phi(\hat{\theta}) \leq \bar{\eta}_k$ for all $t \geq T_k$, allowing us to apply Lemma 4 with smaller and smaller values of $\bar{\eta}$ until it reduces to the statistical error $\bar{\epsilon}_{\text{stat}}$.

At the first iteration, we have no a priori bound on the error $\bar{\eta}_0 = \phi(\theta^0) - \phi(\hat{\theta})$. However, since Lemma 4 involves the quantity $\varepsilon = \min(\bar{\eta}/\lambda_n, \bar{\rho})$, we may still apply it⁶ at the first epoch with $\varepsilon_0 = \bar{\rho}$ and $T_0 = 0$. In this way, we conclude that for all $t \geq 0$,

$$\phi(\theta^t) - \phi(\hat{\theta}) \leq \kappa^t(\phi(\theta^0) - \phi(\hat{\theta})) + \frac{2}{1-\kappa}\xi\beta(\bar{\rho}^2 + \bar{\epsilon}_{\text{stat}}^2).$$

Now since the contraction coefficient $\kappa \in (0, 1)$, for all iterations $t \geq T_1 := (\lceil \log(2\bar{\eta}_0/\bar{\eta}_1) / \log(1/\kappa) \rceil)_+$, we are guaranteed that

$$\phi(\theta^t) - \phi(\hat{\theta}) \leq \underbrace{\frac{4\xi\beta}{1-\kappa}(\bar{\rho}^2 + \bar{\epsilon}_{\text{stat}}^2)}_{\bar{\eta}_1} \leq \frac{8\xi\beta}{1-\kappa} \max(\bar{\rho}^2, \bar{\epsilon}_{\text{stat}}^2).$$

This same argument can now be applied in a recursive manner. Suppose that for some $k \geq 1$, we are given a pair $(\bar{\eta}_k, T_k)$ such that condition (56) holds. An application of Lemma 4 yields the bound

$$\phi(\theta^t) - \phi(\hat{\theta}) \leq \kappa^{t-T_k}(\phi(\theta^{T_k}) - \phi(\hat{\theta})) + \frac{2\xi\beta}{1-\kappa}(\varepsilon_k^2 + \bar{\epsilon}_{\text{stat}}^2) \quad \text{for all } t \geq T_k.$$

We now define $\bar{\eta}_{k+1} := \frac{4\xi\beta}{1-\kappa}(\varepsilon_k^2 + \bar{\epsilon}_{\text{stat}}^2)$. Once again, since $\kappa < 1$ by assumption, we can choose $T_{k+1} := \lceil \log(2\bar{\eta}_k/\bar{\eta}_{k+1}) / \log(1/\kappa) \rceil + T_k$, thereby ensuring that for all $t \geq T_{k+1}$, we have

$$\phi(\theta^t) - \phi(\hat{\theta}) \leq \frac{8\xi\beta}{1-\kappa} \max(\varepsilon_k^2, \bar{\epsilon}_{\text{stat}}^2).$$

⁶It is for precisely this reason that our regularized M -estimator includes the additional side-constraint defined in terms of $\bar{\rho}$.

In this way, we arrive at recursive inequalities involving the tolerances $\{\bar{\eta}_k\}_{k=0}^\infty$ and time steps $\{T_k\}_{k=0}^\infty$ —namely

$$\bar{\eta}_{k+1} \leq \frac{8\xi\beta}{1-\kappa} \max(\varepsilon_k^2, \bar{\epsilon}_{\text{stat}}^2), \quad \text{where } \varepsilon_k = 2 \min\{\bar{\eta}_k/\lambda_n, \bar{\rho}\}, \text{ and} \quad (58a)$$

$$T_k \leq k + \frac{\log(2^k \bar{\eta}_0/\bar{\eta}_k)}{\log(1/\kappa)}. \quad (58b)$$

Now we claim that the recursion (58a) can be unwrapped so as to show that

$$\bar{\eta}_{k+1} \leq \frac{\bar{\eta}_k}{4^{2^{k-1}}} \quad \text{and} \quad \frac{\bar{\eta}_{k+1}}{\lambda_n} \leq \frac{\bar{\rho}}{4^{2^k}} \quad \text{for all } k = 1, 2, \dots \quad (59)$$

Taking these statements as given for the moment, let us now show how they can be used to upper bound the smallest k such that $\bar{\eta}_k \leq \delta^2$. If we are in the first epoch, the claim of the theorem is straightforward from equation (58a). If not, we first use the recursion (59) to upper bound the number of epochs needed and then use the inequality (58b) to obtain the stated result on the total number of iterations needed. Using the second inequality in the recursion (59), we see that it is sufficient to ensure that $\frac{\bar{\rho}\lambda_n}{4^{2^{k-1}}} \leq \delta^2$. Rearranging this inequality, we find that the error drops below δ^2 after at most

$$k_\delta \geq \log \left(\log \left(\frac{\bar{\rho}\lambda_n}{\delta^2} \right) / \log(4) \right) / \log(2) + 1 = \log_2 \log_2 \left(\frac{\bar{\rho}\lambda_n}{\delta^2} \right)$$

epochs. Combining the above bound on k_δ with the recursion 58b, we conclude that the inequality $\phi(\theta^t) - \phi(\hat{\theta}) \leq \delta^2$ is guaranteed to hold for all iterations

$$t \geq k_\delta \left(1 + \frac{\log 2}{\log(1/\kappa)} \right) + \frac{\log \frac{\bar{\eta}_0}{\delta^2}}{\log(1/\kappa)},$$

which is the desired result.

It remains to prove the recursion (59), which we do via induction on the index k . We begin with base case $k = 1$. Recalling the setting of $\bar{\eta}_1$ and our assumption on λ_n in the theorem statement (30), we are guaranteed that $\bar{\eta}_1/\lambda_n \leq \bar{\rho}/4$, so that $\varepsilon_1 \leq \varepsilon_0 = \bar{\rho}$. By applying equation (58a) with $\varepsilon_1 = 2\bar{\eta}_1/\lambda_n$ and assuming $\varepsilon_1 \geq \bar{\epsilon}_{\text{stat}}$, we obtain

$$\bar{\eta}_2 \leq \frac{32\xi\beta\bar{\eta}_1^2}{(1-\kappa)\lambda_n^2} \stackrel{(i)}{\leq} \frac{32\xi\beta\bar{\rho}\bar{\eta}_1}{(1-\kappa)4\lambda_n} \stackrel{(ii)}{\leq} \frac{\bar{\eta}_1}{4}, \quad (60)$$

where step (i) uses the fact that $\frac{\bar{\eta}_1}{\lambda_n} \leq \frac{\bar{\rho}}{4}$, and step (ii) uses the condition (30) on λ_n . We have thus verified the first inequality (59) for $k = 1$. Turning to the second inequality in the statement (59), using equation 60, we have

$$\frac{\bar{\eta}_2}{\lambda_n} \leq \frac{\bar{\eta}_1}{4\lambda_n} \stackrel{(iii)}{\leq} \frac{\bar{\rho}}{16},$$

where step (iii) follows from the assumption (30) on λ_n . Turning to the inductive step, we again assume that $2\bar{\eta}_k/\lambda_n \geq \bar{\epsilon}_{\text{stat}}$ and obtain from inequality (58a)

$$\bar{\eta}_{k+1} \leq \frac{32\xi\beta\bar{\eta}_k^2}{(1-\kappa)\lambda_n^2} \stackrel{(iv)}{\leq} \frac{32\xi\beta\bar{\eta}_k\bar{\rho}}{(1-\kappa)\lambda_n4^{2^{k-1}}} \stackrel{(v)}{\leq} \frac{\bar{\eta}_k}{4^{2^{k-1}}}.$$

Here step (iv) uses the second inequality of the inductive hypothesis (59) and step (v) is a consequence of the condition on λ_n as before. The second part of the induction is similarly established, completing the proof.

5.3 Proof of Corollary 1

In order to prove this claim, we must show that $\epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}})$, as defined in equation (23), is of order lower than $\mathbb{E}[\|\widehat{\theta} - \theta^*\|^2] = \mathbb{E}[\|\Delta^*\|^2]$. We make use of the following lemma, proved in Appendix C:

Lemma 5. *If $\rho \leq \mathcal{R}(\theta^*)$, then for any solution $\widehat{\theta}$ of the constrained problem (1) and any \mathcal{R} -decomposable subspace pair $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$, the statistical error $\Delta^* = \widehat{\theta} - \theta^*$ satisfies the inequality*

$$\mathcal{R}(\Delta^*) \leq 2\Psi(\overline{\mathcal{M}}^\perp)\|\Delta^*\| + \mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)). \quad (61)$$

Using this lemma, we can complete the proof of Corollary 1. Recalling the form (23), under the condition $\theta^* \in \mathcal{M}$, we have

$$\epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}}) := \frac{32(\tau_u(\mathcal{L}_n) + \tau_\ell(\mathcal{L}_n)) (2\mathcal{R}(\Delta^*) + \Psi(\overline{\mathcal{M}}^\perp)\|\Delta^*\|)^2}{\gamma_u}.$$

Using the assumption $\frac{(\tau_u(\mathcal{L}_n) + \tau_\ell(\mathcal{L}_n))\Psi^2(\overline{\mathcal{M}}^\perp)}{\gamma_u} = o(1)$, it suffices to show that $\mathcal{R}(\Delta^*) \leq 2\Psi(\overline{\mathcal{M}}^\perp)\|\Delta^*\|$. Since Corollary 1 assumes that $\theta^* \in \mathcal{M}$ and hence that $\Pi_{\mathcal{M}^\perp}(\theta^*) = 0$, Lemma 5 implies that $\mathcal{R}(\Delta^*) \leq 2\Psi(\overline{\mathcal{M}}^\perp)\|\Delta^*\|$, as required.

5.4 Proofs of Corollaries 2 and 3

The central challenge in proving this result is verifying that suitable forms of the RSC and RSM conditions hold with sufficiently small parameters $\tau_\ell(\mathcal{L}_n)$ and $\tau_u(\mathcal{L}_n)$.

Lemma 6. *Define the maximum variance $\zeta(\Sigma) := \max_{j=1,2,\dots,d} \Sigma_{jj}$. Under the conditions of Corollary 2, there are universal positive constants (c_0, c_1) such that for all $\Delta \in \mathbb{R}^d$, we have*

$$\frac{\|X\Delta\|_2^2}{n} \geq \frac{1}{2}\|\Sigma^{1/2}\Delta\|_2^2 - c_1\zeta(\Sigma)\frac{\log d}{n}\|\Delta\|_1^2, \quad \text{and} \quad (62a)$$

$$\frac{\|X\Delta\|_2^2}{n} \leq 2\|\Sigma^{1/2}\Delta\|_2^2 + c_1\zeta(\Sigma)\frac{\log d}{n}\|\Delta\|_1^2, \quad (62b)$$

with probability at least $1 - \exp(-c_0 n)$.

Note that this lemma implies that the RSC and RSM conditions both hold with high probability, in particular with parameters

$$\begin{aligned} \gamma_\ell &= \frac{1}{2}\sigma_{\min}(\Sigma), \text{ and} & \tau_\ell(\mathcal{L}_n) &= c_1\zeta(\Sigma)\frac{\log d}{n}, & \text{for RSC, and} \\ \gamma_u &= 2\sigma_{\max}(\Sigma) \text{ and} & \tau_u(\mathcal{L}_n) &= c_1\zeta(\Sigma)\frac{\log d}{n} & \text{for RSM.} \end{aligned}$$

This lemma has been proved by Raskutti et al. [34] for obtaining minimax rates in sparse linear regression.

Let us first prove Corollary 2 in the special case of hard sparsity ($q = 0$), in which θ^* is supported on a subset S of cardinality s . Let us define the model subspace $\mathcal{M} := \{\theta \in \mathbb{R}^d \mid \theta_j = 0 \text{ for all } j \notin S\}$, so that $\theta^* \in \mathcal{M}$. Recall from Section 2.4.1 that the ℓ_1 -norm is decomposable with respect to \mathcal{M} and \mathcal{M}^\perp ; as a consequence, we may also set $\overline{\mathcal{M}}^\perp = \mathcal{M}$ in the definitions (22) and (23). By definition (12) of the subspace compatibility between with ℓ_1 -norm as the regularizer, and ℓ_2 -norm as the error norm, we have $\Psi^2(\mathcal{M}) = s$. Using the settings of $\tau_\ell(\mathcal{L}_n)$ and $\tau_u(\mathcal{L}_n)$ guaranteed by Lemma 6 and substituting into equation (22), we obtain a contraction coefficient

$$\kappa(\Sigma) := \left\{ 1 - \frac{\sigma_{\min}(\Sigma)}{4\sigma_{\max}(\Sigma)} + \chi_n(\Sigma) \right\} \left\{ 1 - \chi_n(\Sigma) \right\}^{-1}, \quad (63)$$

where $\chi_n(\Sigma) := \frac{c_2 \zeta(\Sigma)}{\sigma_{\max}(\Sigma)} \frac{s \log d}{n}$ for some universal constant c_2 . A similar calculation shows that the tolerance term takes the form

$$\epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}}) \leq c_3 \chi_n(\Sigma) \left\{ \frac{\|\Delta^*\|_1^2}{s} + \|\Delta^*\|_2^2 \right\} \quad \text{for some constant } c_3.$$

Since $\rho \leq \|\theta^*\|_1$, then Lemma 5 (as exploited in the proof of Corollary 1) shows that $\|\Delta^*\|_1^2 \leq 4s\|\Delta^*\|_2^2$, and hence that $\epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}}) \leq c_3 \chi_n(\Sigma) \|\Delta^*\|_2^2$. This completes the proof of the claim (36) for $q = 0$.

We now turn to the case $q \in (0, 1]$, for which we bound the term $\epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}})$ using a slightly different choice of the subspace pair \mathcal{M} and $\overline{\mathcal{M}}^\perp$. For a truncation level $\mu > 0$ to be chosen, define the set $S_\mu := \{j \in \{1, 2, \dots, d\} \mid |\theta_j^*| > \mu\}$, and define the associated subspaces $\mathcal{M} = \mathcal{M}(S_\mu)$ and $\overline{\mathcal{M}}^\perp = \mathcal{M}^\perp(S_\mu)$. By combining Lemma 5 and the definition (23) of $\epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}})$, for any pair $(\mathcal{M}(S_\mu), \mathcal{M}^\perp(S_\mu))$, we have

$$\epsilon^2(\Delta^*; \mathcal{M}, \mathcal{M}^\perp) \leq \frac{c \zeta(\Sigma)}{\sigma_{\max}(\Sigma)} \frac{\log d}{n} (\|\Pi_{\mathcal{M}^\perp}(\theta^*)\|_1 + \sqrt{|S_\mu|} \|\Delta^*\|_2)^2,$$

where to simplify notation, we have omitted the dependence of \mathcal{M} and \mathcal{M}^\perp on S_μ . We now choose the threshold μ optimally, so as to trade-off the term $\|\Pi_{\mathcal{M}^\perp}(\theta^*)\|_1$, which decreases as μ increases, with the term $\sqrt{|S_\mu|} \|\Delta^*\|_2$, which increases as μ increases.

By definition of $\mathcal{M}^\perp(S_\mu)$, we have

$$\|\Pi_{\mathcal{M}^\perp}(\theta^*)\|_1 = \sum_{j \notin S_\mu} |\theta_j^*| = \mu \sum_{j \notin S_\mu} \frac{|\theta_j^*|}{\mu} \leq \mu \sum_{j \notin S_\mu} \left(\frac{|\theta_j^*|}{\mu} \right)^q,$$

where the inequality holds since $|\theta_j^*| \leq \mu$ for all $j \notin S_\mu$. Now since $\theta^* \in \mathbb{B}_q(R_q)$, we conclude that

$$\|\Pi_{\mathcal{M}^\perp}(\theta^*)\|_1 \leq \mu^{1-q} \sum_{j \notin S_\mu} |\theta_j^*|^q \leq \mu^{1-q} R_q. \quad (64)$$

On the other hand, again using the inclusion $\theta^* \in \mathbb{B}_q(R_q)$, we have $R_q \geq \sum_{j \in S_\mu} |\theta_j^*|^q \geq |S_\mu| \mu^q$ which implies that $|S_\mu| \leq \mu^{-q} R_q$. By combining this bound with inequality (64), we obtain the upper bound

$$\epsilon^2(\Delta^*; \mathcal{M}, \mathcal{M}^\perp) \leq \frac{c \zeta(\Sigma)}{\sigma_{\max}(\Sigma)} \frac{\log d}{n} (\mu^{2-2q} R_q^2 + \mu^{-q} R_q \|\Delta^*\|_2^2) = \frac{c \zeta(\Sigma)}{\sigma_{\max}(\Sigma)} \frac{\log d}{n} \mu^{-q} R_q (\mu^{2-q} R_q + \|\Delta^*\|_2^2).$$

Setting $\mu^2 = \frac{\log d}{n}$ then yields

$$\epsilon^2(\Delta^*; \mathcal{M}, \mathcal{M}^\perp) \leq \chi_n(\Sigma) \left\{ R_q \left(\frac{\log d}{n} \right)^{1-q/2} + \|\Delta^*\|_2^2 \right\}, \quad \text{where } \chi_n(\Sigma) := \frac{c\zeta(\Sigma)}{\sigma_{\max}(\Sigma)} R_q \left(\frac{\log d}{n} \right)^{1-q/2}.$$

Finally, let us verify the stated form of the contraction coefficient. For the given subspace $\overline{\mathcal{M}}^\perp = \mathcal{M}(S_\mu)$ and choice of μ , we have $\Psi^2(\overline{\mathcal{M}}^\perp) = |S_\mu| \leq \mu^{-q} R_q$. From Lemma 6, we have

$$16\Psi^2(\overline{\mathcal{M}}^\perp) \frac{\tau_\ell(\mathcal{L}_n) + \tau_u(\mathcal{L}_n)}{\gamma_u} \leq \chi_n(\Sigma),$$

and hence, by definition (22) of the contraction coefficient,

$$\kappa \leq \left\{ 1 - \frac{\gamma_\ell}{2\gamma_u} + \chi_n(\Sigma) \right\} \left\{ 1 - \chi_n(\Sigma) \right\}^{-1}.$$

For proving Corollary 3, we observe that the stated settings $\overline{\gamma}_\ell$, $\chi_n(\Sigma)$ and κ follow directly from Lemma 6. The bound for condition 2(a) follows from a standard argument about the suprema of d independent Gaussians with variance ν .

5.5 Proof of Corollary 4

This proof is analogous to that of Corollary 2, but appropriately adapted to the matrix setting. We first state a lemma that allows us to establish appropriate forms of the RSC/RSM conditions. Recall that we are studying an instance of matrix regression with random design, where the vectorized form $\text{vec}(X)$ of each matrix is drawn from a $N(0, \Sigma)$ distribution, where $\Sigma \in \mathbb{R}^{d^2 \times d^2}$ is some covariance matrix. In order to state this result, let us define the quantity

$$\zeta_{\text{mat}}(\Sigma) := \sup_{\|u\|_2=1, \|v\|_2=1} \text{var}(u^T X v), \quad \text{where } \text{vec}(X) \sim N(0, \Sigma). \quad (65)$$

Lemma 7. *Under the conditions of Corollary 4, there are universal positive constants (c_0, c_1) such that*

$$\frac{\|\mathfrak{X}_n(\Delta)\|_2^2}{n} \geq \frac{1}{2} \sigma_{\min}(\Sigma) \|\Delta\|_F^2 - c_1 \zeta_{\text{mat}}(\Sigma) \frac{d}{n} \|\Delta\|_1^2, \quad \text{and} \quad (66a)$$

$$\frac{\|\mathfrak{X}_n(\Delta)\|_2^2}{n} \leq 2\sigma_{\max}(\Sigma) \|\Delta\|_F^2 - c_1 \zeta_{\text{mat}}(\Sigma) \frac{d}{n} \|\Delta\|_1^2, \quad \text{for all } \Delta \in \mathbb{R}^{d \times d}. \quad (66b)$$

with probability at least $1 - \exp(-c_0 n)$.

Given the quadratic nature of the least-squares loss, the bound (66a) implies that the RSC condition holds with $\gamma_\ell = \frac{1}{2}\sigma_{\min}(\Sigma)$ and $\tau_\ell(\mathcal{L}_n) = c_1 \zeta_{\text{mat}}(\Sigma) \frac{d}{n}$, whereas the bound (66b) implies that the RSM condition holds with $\gamma_u = 2\sigma_{\max}(\Sigma)$ and $\tau_u(\mathcal{L}_n) = c_1 \zeta_{\text{mat}}(\Sigma) \frac{d}{n}$.

We now prove Corollary 4 in the special case of exactly low rank matrices ($q = 0$), in which Θ^* has some rank $r \leq d$. Given the singular value decomposition $\Theta^* = UDV^T$, let U^r and V^r be the $d \times r$ matrices whose columns correspond to the r non-zero (left and right, respectively) singular vectors of Θ^* . As in Section 2.4.2, define the subspace of matrices

$$\mathcal{M}(U^r, V^r) := \{ \Theta \in \mathbb{R}^{d \times d} \mid \text{col}(\Theta) \subseteq U^r \text{ and } \text{row}(\Theta) \subseteq V^r \}, \quad (67)$$

as well as the associated set $\overline{\mathcal{M}}^\perp(U^r, V^r)$. Note that $\Theta^* \in \mathcal{M}$ by construction, and moreover (as discussed in Section 2.4.2, the nuclear norm is decomposable with respect to the pair $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$).

By definition (12) of the subspace compatibility with nuclear norm as the regularizer and Frobenius norm as the error norm, we have $\Psi^2(\mathcal{M}) = r$. Using the settings of $\tau_\ell(\mathcal{L}_n)$ and $\tau_u(\mathcal{L}_n)$ guaranteed by Lemma 7 and substituting into equation (22), we obtain a contraction coefficient

$$\kappa(\Sigma) := \left\{ 1 - \frac{\sigma_{\min}(\Sigma)}{4\sigma_{\max}(\Sigma)} + \chi_n(\Sigma) \right\} \left\{ 1 - \chi_n(\Sigma) \right\}^{-1}, \quad (68)$$

where $\chi_n(\Sigma) := \frac{c_2 \zeta_{\text{mat}}(\Sigma)}{\sigma_{\max}(\Sigma)} \frac{rd}{n}$ for some universal constant c_2 . A similar calculation shows that the tolerance term takes the form

$$\epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}}) \leq c_3 \chi_n(\Sigma) \left\{ \frac{\|\Delta^*\|_1^2}{r} + \|\Delta^*\|_F^2 \right\} \quad \text{for some constant } c_3.$$

Since $\rho \leq \|\Theta^*\|_1$ by assumption, Lemma 5 (as exploited in the proof of Corollary 1) shows that $\|\Delta^*\|_1^2 \leq 4r\|\Delta^*\|_F^2$, and hence that

$$\epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}}) \leq c_3 \chi_n(\Sigma) \|\Delta^*\|_F^2,$$

which show the claim (42) for $q = 0$.

We now turn to the case $q \in (0, 1]$; as in the proof of this case for Corollary 2, we bound $\epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}})$ using a slightly different choice of the subspace pair. Recall our notation $\sigma_1(\Theta^*) \geq \sigma_2(\Theta^*) \geq \dots \geq \sigma_d(\Theta^*) \geq 0$ for the ordered singular values of Θ^* . For a threshold μ to be chosen, define $S_\mu = \{j \in \{1, 2, \dots, d\} \mid \sigma_j(\Theta^*) > \mu\}$, and $U(S_\mu) \in \mathbb{R}^{d \times |S_\mu|}$ be the matrix of left singular vectors indexed by S_μ , with the matrix $V(S_\mu)$ defined similarly. We then define the subspace $\mathcal{M}(S_\mu) := \mathcal{M}(U(S_\mu), V(S_\mu))$ in an analogous fashion to equation (67), as well as the subspace $\overline{\mathcal{M}}^\perp(S_\mu)$.

Now by a combination of Lemma 5 and the definition (23) of $\epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}})$, for any pair $(\mathcal{M}(S_\mu), \overline{\mathcal{M}}^\perp(S_\mu))$, we have

$$\epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}}^\perp) \leq \frac{c \zeta_{\text{mat}}(\Sigma)}{\sigma_{\max}(\Sigma)} \frac{d}{n} \left(\sum_{j \notin S_\mu} \sigma_j(\Theta^*) + \sqrt{|S_\mu|} \|\Delta^*\|_F \right)^2,$$

where to simplify notation, we have omitted the dependence of \mathcal{M} and \mathcal{M}^\perp on S_μ . As in the proof of Corollary 2, we now choose the threshold μ optimally, so as to trade-off the term $\sum_{j \notin S_\mu} \sigma_j(\Theta^*)$ with its competitor $\sqrt{|S_\mu|} \|\Delta^*\|_F$. Exploiting the fact that $\Theta^* \in \mathbb{B}_q(R_q)$ and following the same steps as the proof of Corollary 2 yields the bound

$$\epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}}^\perp) \leq \frac{c \zeta_{\text{mat}}(\Sigma)}{\sigma_{\max}(\Sigma)} \frac{d}{n} (\mu^{2-2q} R_q^2 + \mu^{-q} R_q \|\Delta^*\|_F^2).$$

Setting $\mu^2 = \frac{d}{n}$ then yields

$$\epsilon^2(\Delta^*; \mathcal{M}, \overline{\mathcal{M}}^\perp) \leq \chi_n(\Sigma) \left\{ R_q \left(\frac{d}{n} \right)^{1-q/2} + \|\Delta^*\|_F^2 \right\},$$

as claimed. The stated form of the contraction coefficient can be verified by a calculation analogous to the proof of Corollary 2.

5.6 Proof of Corollary 5

In this case, we let $\mathfrak{X}_n : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^n$ be the operator defined by the model of random signed matrix sampling [30]. As previously argued, establishing the RSM/RSC property amounts to obtaining a form of uniform control over $\frac{\|\mathfrak{X}_n(\Theta)\|_2^2}{n}$. More specifically, from the proof of Theorem 1, we see that it suffices to have a form of RSC for the difference $\widehat{\Delta}^t = \Theta^t - \widehat{\Theta}$, and a form of RSM for the difference $\Theta^{t+1} - \Theta^t$. The following two lemmas summarize these claims:

Lemma 8. *There is a constant c such that for all iterations $t = 0, 1, 2, \dots$ and integers $r = 1, 2, \dots, d-1$, with probability at least $1 - \exp(-d \log d)$,*

$$\frac{\|\mathfrak{X}_n(\widehat{\Delta}^t)\|_2^2}{n} \geq \frac{1}{2} \|\widehat{\Delta}^t\|_F^2 - \underbrace{c\alpha \sqrt{\frac{rd \log d}{n}} \left\{ \frac{\sum_{j=r+1}^d \sigma_j(\Theta^*)}{\sqrt{r}} + \alpha \sqrt{\frac{rd \log d}{n}} + \|\Delta^*\|_F \right\}}_{\delta_\ell(r)}. \quad (69)$$

Lemma 9. *There is a constant c such that for all iterations $t = 0, 1, 2, \dots$ and integers $r = 1, 2, \dots, d-1$, with probability at least $1 - \exp(-d \log d)$, the difference $\Gamma^t := \Theta^{t+1} - \Theta^t$ satisfies the inequality $\frac{\|\mathfrak{X}_n(\Gamma^t)\|_2^2}{n} \leq 2\|\Gamma^t\|_F^2 + \delta_u(r)$, where*

$$\delta_u(r) := c\alpha \sqrt{\frac{rd \log d}{n}} \left\{ \frac{\sum_{j=r+1}^d \sigma_j(\Theta^*)}{\sqrt{r}} + \alpha \sqrt{\frac{rd \log d}{n}} + \|\Delta^*\|_F + \|\widehat{\Delta}^t\|_F + \|\widehat{\Delta}^{t+1}\|_F \right\}.$$

We can now complete the proof of Corollary 5 by a minor modification of the proof of Theorem 1. Recalling the elementary relation (53), we have

$$\|\Theta^{t+1} - \widehat{\Theta}\|_F^2 = \|\Theta^t - \widehat{\Theta}\|_F^2 + \|\Theta^t - \Theta^{t+1}\|_F^2 - 2\langle\langle \Theta^t - \widehat{\Theta}, \Theta^t - \Theta^{t+1} \rangle\rangle.$$

From the proof of Lemma 2, we see that the combination of Lemma 8 and 9 (with $\gamma_\ell = \frac{1}{2}$ and $\gamma_u = 2$) imply that

$$2\langle\langle \Theta^t - \Theta^{t+1}, \Theta^t - \widehat{\Theta} \rangle\rangle \geq \|\Theta^t - \Theta^{t+1}\|_F^2 + \frac{1}{4} \|\Theta^t - \widehat{\Theta}\|_F^2 - \delta_u(r) - \delta_\ell(r)$$

and hence that

$$\|\widehat{\Delta}^{t+1}\|_F^2 \leq \frac{3}{4} \|\widehat{\Delta}^t\|_F^2 + \delta_\ell(r) + \delta_u(r).$$

We substitute the forms of $\delta_\ell(r)$ and $\delta_u(r)$ given in Lemmas 8 and 9 respectively; performing some algebra then yields

$$\left\{ 1 - \frac{c\alpha \sqrt{\frac{rd \log d}{n}}}{\|\widehat{\Delta}^{t+1}\|_F} \right\} \|\widehat{\Delta}^{t+1}\|_F^2 \leq \left\{ \frac{3}{4} + \frac{c\alpha \sqrt{\frac{rd \log d}{n}}}{\|\widehat{\Delta}^t\|_F} \right\} \|\widehat{\Delta}^t\|_F^2 + c' \delta_\ell(r).$$

Consequently, as long as $\min\{\|\widehat{\Delta}^t\|_F^2, \|\widehat{\Delta}^{t+1}\|_F^2\} \geq c_3 \alpha \frac{rd \log d}{n}$ for a sufficiently large constant c_3 , we are guaranteed the existence of some $\kappa_t \in (0, 1)$ decreasing with t such that

$$\|\widehat{\Delta}^{t+1}\|_F^2 \leq \kappa \|\widehat{\Delta}^t\|_F^2 + c' \delta_\ell(r). \quad (70)$$

Since $\delta_\ell(r) = \Omega(\frac{rd \log d}{n})$, this inequality (70) is valid for all $t = 0, 1, 2, \dots$ as long as c' is sufficiently large. Now iterating this bound, we see that

$$\|\widehat{\Delta}^{t+1}\|_F^2 \leq \left(\prod_{s=1}^t \kappa_s \right) \|\widehat{\Delta}^0\|_F^2 + c' \delta_\ell(r) \left(\kappa_t + \kappa_t \kappa_{t-1} + \dots + \prod_{s=2}^t \kappa_s \right).$$

Since κ_t is decreasing in t , we observe that the second term in the above bound is at most

$$c' \delta_\ell(r) \left(\kappa_t + \kappa_t \kappa_{t-1} + \dots + \prod_{s=2}^t \kappa_s \right) \leq c' \delta_\ell(r) \left(\kappa_1 + \kappa_1^2 + \kappa_1^{t-1} \right) \leq c' \frac{\delta_\ell(r)}{1 - \kappa_1}.$$

We also define $\bar{\kappa}_t = (\sum_{s=1}^t \kappa_s)/t$. Then the arithmetic mean-geometric mean inequality yields the upper bound $\prod_{s=1}^t \kappa_s \leq \bar{\kappa}_t^t$. Combining this with our earlier upper bound further yields the inequality

$$\|\widehat{\Delta}^{t+1}\|_F^2 \leq \bar{\kappa}_t^t \|\widehat{\Delta}^0\|_F^2 + \frac{c'}{1 - \kappa_1} \delta_\ell(r). \quad (71)$$

It remains to choose the cut-off $r \in \{1, 2, \dots, d-1\}$ so as to minimize the term $\delta_\ell(r)$. In particular, when $\Theta^* \in \mathbb{B}_q(R_q)$, then as shown in the paper [29], the optimal choice is $r \asymp \alpha^{-q} R_q \left(\frac{n}{d \log d} \right)^{q/2}$. Substituting into the inequality (71) and performing some algebra yields that there is a universal constant c_4 such that the bound

$$\|\widehat{\Delta}^{t+1}\|_F^2 \leq \kappa^t \|\widehat{\Delta}^0\|_F^2 + \frac{c_4}{1 - \kappa} \left\{ R_q \left(\frac{\alpha d \log d}{n} \right)^{1-q/2} + \sqrt{R_q \left(\frac{\alpha d \log d}{n} \right)^{1-q/2}} \|\Delta^*\|_F \right\}.$$

holds. Now by the Cauchy-Schwarz inequality we have

$$\sqrt{R_q \left(\frac{\alpha d \log d}{n} \right)^{1-q/2}} \|\Delta^*\|_F \leq \frac{1}{2} R_q \left(\frac{\alpha d \log d}{n} \right)^{1-q/2} + \frac{1}{2} \|\Delta^*\|_F^2,$$

and the claimed inequality (45) follows.

5.7 Proof of Corollary 6

Again the main argument in the proof would be to establish the RSM and RSC properties for the decomposition problem. We define $\widehat{\Delta}_\Theta^t = \Theta^t - \widehat{\Theta}$ and $\widehat{\Delta}_\Gamma^t = \Gamma^t - \widehat{\Gamma}$. We start with giving a lemma that establishes RSC for the differences $(\widehat{\Delta}_\Theta^t, \widehat{\Delta}_\Gamma^t)$. We recall that just like noted in the previous section, it suffices to show RSC only for these differences. Showing RSC/RSM in this example amounts to analyzing $\|\widehat{\Delta}_\Theta^t + \widehat{\Delta}_\Gamma^t\|_F^2$. We recall that this section assumes that Γ^* has only s non-zero columns.

Lemma 10. *There is a constant c such that for all iterations $t = 0, 1, 2, \dots$,*

$$\|\widehat{\Delta}_\Theta^t + \widehat{\Delta}_\Gamma^t\|_F^2 \geq \frac{1}{2} (\|\widehat{\Delta}_\Theta^t\|_F^2 + \|\widehat{\Delta}_\Gamma^t\|_F^2) - c\alpha \sqrt{\frac{s}{d_2}} \left(\|\widehat{\Gamma} - \Gamma^*\|_F + \alpha \sqrt{\frac{s}{d_2}} \right) \quad (72)$$

This proof of this lemma follows by a straightforward modification of analogous results in the paper [1].

Matrix decomposition has the interesting property that the RSC condition holds in a deterministic sense (as opposed to with high probability). The same deterministic guarantee holds for the RSM condition; indeed, we have

$$\|\widehat{\Delta}_\Delta^t + \widehat{\Delta}_\Gamma^t\|_F^2 \leq 2(\|\widehat{\Delta}_\Theta^t\|_F^2 + \|\widehat{\Delta}_\Gamma^t\|_F^2), \quad (73)$$

by Cauchy-Schwartz inequality. Now we appeal to the more general form of Theorem 1 as stated in Equation 49, which gives

$$\|\widehat{\Delta}_\Theta^{t+1}\|_F^2 + \|\widehat{\Delta}_\Gamma^{t+1}\|_F^2 \leq \left(\frac{3}{4}\right)^t (\|\widehat{\Delta}_\Theta^0\|_F^2 + \|\widehat{\Delta}_\Gamma^0\|_F^2) + c\sqrt{\frac{\alpha s}{d_2}} \left(\|\widehat{\Gamma} - \Gamma^*\|_F + \frac{\alpha s}{d_2}\right).$$

The stated form of the corollary follows by an application of Cauchy-Schwarz inequality.

6 Discussion

In this paper, we have shown that even though high-dimensional M -estimators in statistics are neither strongly convex nor smooth, simple first-order methods can still enjoy global guarantees of geometric convergence. The key insight is that strong convexity and smoothness need only hold in restricted senses, and moreover, these conditions are satisfied with high probability for many statistical models and decomposable regularizers used in practice. Examples include sparse linear regression and ℓ_1 -regularization, various statistical models with group-sparse regularization, matrix regression with nuclear norm constraints (including matrix completion and multi-task learning), and matrix decomposition problems. Overall, our results highlight some important connections between computation and statistics: the properties of M -estimators favorable for fast rates in a statistical sense can also be used to establish fast rates for optimization algorithms.

Acknowledgements: All three authors were partially supported by grants AFOSR-09NL184; in addition, AA was partially supported by a Microsoft Graduate Fellowship and Google PhD Fellowship, and SN and MJW acknowledge funding from NSF-CDI-0941742. We would like to thank the anonymous reviewers and associate editor for their helpful comments that helped to improve the paper, and Bin Yu for inspiring discussions on the interaction between statistical and optimization error.

A Auxiliary results for Theorem 1

In this appendix, we provide the proofs of various auxiliary lemmas required in the proof of Theorem 1.

A.1 Proof of Lemma 1

Since θ^t and $\widehat{\theta}$ are both feasible and $\widehat{\theta}$ lies on the constraint boundary, we have $\mathcal{R}(\theta^t) \leq \mathcal{R}(\widehat{\theta})$. Since $\mathcal{R}(\widehat{\theta}) \leq \mathcal{R}(\theta^*) + \mathcal{R}(\widehat{\theta} - \theta^*)$ by triangle inequality, we conclude that

$$\mathcal{R}(\theta^t) \leq \mathcal{R}(\theta^*) + \mathcal{R}(\Delta^*).$$

Since $\theta^* = \Pi_{\mathcal{M}}(\theta^*) + \Pi_{\mathcal{M}^\perp}(\theta^*)$, a second application of triangle inequality yields

$$\mathcal{R}(\theta^t) \leq \mathcal{R}(\Pi_{\mathcal{M}}(\theta^*)) + \mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)) + \mathcal{R}(\Delta^*). \quad (74)$$

Now define the difference $\Delta^t := \theta^t - \theta^*$. (Note that this is slightly different from $\widehat{\Delta}^t$, which is measured relative to the optimum $\widehat{\theta}$.) With this notation, we have

$$\begin{aligned} \mathcal{R}(\theta^t) &= \mathcal{R}(\Pi_{\mathcal{M}}(\theta^*) + \Pi_{\mathcal{M}^\perp}(\theta^*) + \Pi_{\bar{\mathcal{M}}}(\Delta^t) + \Pi_{\bar{\mathcal{M}}^\perp}(\Delta^t)) \\ &\stackrel{(i)}{\geq} \mathcal{R}(\Pi_{\mathcal{M}}(\theta^*) + \Pi_{\bar{\mathcal{M}}^\perp}(\Delta^t)) - \mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*) + \Pi_{\bar{\mathcal{M}}}(\Delta^t)) \\ &\stackrel{(ii)}{\geq} \mathcal{R}(\Pi_{\mathcal{M}}(\theta^*) + \Pi_{\bar{\mathcal{M}}^\perp}(\Delta^t)) - \mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)) - \mathcal{R}(\Pi_{\bar{\mathcal{M}}}(\Delta^t)), \end{aligned}$$

where steps (i) and (ii) each use the triangle inequality. Now by the decomposability condition, we have $\mathcal{R}(\Pi_{\mathcal{M}}(\theta^*) + \Pi_{\bar{\mathcal{M}}^\perp}(\Delta^t)) = \mathcal{R}(\Pi_{\mathcal{M}}(\theta^*)) + \mathcal{R}(\Pi_{\bar{\mathcal{M}}^\perp}(\Delta^t))$, so that we have shown that

$$\mathcal{R}(\Pi_{\mathcal{M}}(\theta^*)) + \mathcal{R}(\Pi_{\bar{\mathcal{M}}^\perp}(\Delta^t)) - \mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)) - \mathcal{R}(\Pi_{\bar{\mathcal{M}}}(\Delta^t)) \leq \mathcal{R}(\theta^t).$$

Combining this inequality with the earlier bound (74) yields

$$\mathcal{R}(\Pi_{\mathcal{M}}(\theta^*)) + \mathcal{R}(\Pi_{\bar{\mathcal{M}}^\perp}(\Delta^t)) - \mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)) - \mathcal{R}(\Pi_{\bar{\mathcal{M}}}(\Delta^t)) \leq \mathcal{R}(\Pi_{\mathcal{M}}(\theta^*)) + \mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)) + \mathcal{R}(\Delta^*).$$

Re-arranging yields the inequality

$$\mathcal{R}(\Pi_{\bar{\mathcal{M}}^\perp}(\Delta^t)) \leq \mathcal{R}(\Pi_{\bar{\mathcal{M}}}(\Delta^t)) + 2\mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)) + \mathcal{R}(\Delta^*). \quad (75)$$

The final step is to translate this inequality into one that applies to the optimization error $\widehat{\Delta}^t = \theta^t - \widehat{\theta}$. Recalling that $\Delta^* = \widehat{\theta} - \theta^*$, we have $\widehat{\Delta}^t = \Delta^t - \Delta^*$, and hence

$$\mathcal{R}(\widehat{\Delta}^t) \leq \mathcal{R}(\Delta^t) + \mathcal{R}(\Delta^*), \quad \text{by triangle inequality.} \quad (76)$$

In addition, we have

$$\begin{aligned} \mathcal{R}(\Delta^t) &\leq \mathcal{R}(\Pi_{\bar{\mathcal{M}}^\perp}(\Delta^t)) + \mathcal{R}(\Pi_{\bar{\mathcal{M}}}(\Delta^t)) \stackrel{(i)}{\leq} 2\mathcal{R}(\Pi_{\bar{\mathcal{M}}}(\Delta^t)) + 2\mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)) + \mathcal{R}(\Delta^*) \\ &\stackrel{(ii)}{\leq} 2\Psi(\bar{\mathcal{M}}^\perp)\|\Pi_{\bar{\mathcal{M}}}(\Delta^t)\| + 2\mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)) + \mathcal{R}(\Delta^*), \end{aligned}$$

where inequality (i) uses the bound (75), and inequality (ii) uses the definition (12) of the subspace compatibility Ψ . Combining with the inequality (76) yields

$$\mathcal{R}(\widehat{\Delta}^t) \leq 2\Psi(\bar{\mathcal{M}}^\perp)\|\Pi_{\bar{\mathcal{M}}}(\Delta^t)\| + 2\mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)) + 2\mathcal{R}(\Delta^*).$$

Since projection onto a subspace is non-expansive, we have $\|\Pi_{\bar{\mathcal{M}}}(\Delta^t)\| \leq \|\Delta^t\|$, and hence

$$\|\Pi_{\bar{\mathcal{M}}}(\Delta^t)\| \leq \|\widehat{\Delta}^t + \Delta^*\| \leq \|\widehat{\Delta}^t\| + \|\Delta^*\|.$$

Combining the pieces, we obtain the claim (50).

A.2 Proof of Lemma 2

We start by applying the RSC assumption to the pair $\hat{\theta}$ and θ^t , thereby obtaining the lower bound

$$\begin{aligned}\mathcal{L}_n(\hat{\theta}) - \frac{\gamma_\ell}{2}\|\hat{\theta} - \theta^t\|^2 &\geq \mathcal{L}_n(\theta^t) + \langle \nabla \mathcal{L}_n(\theta^t), \hat{\theta} - \theta^t \rangle - \tau_\ell(\mathcal{L}_n)\mathcal{R}^2(\theta^t - \hat{\theta}) \\ &= \mathcal{L}_n(\theta^t) + \langle \nabla \mathcal{L}_n(\theta^t), \theta^{t+1} - \theta^t \rangle + \langle \nabla \mathcal{L}_n(\theta^t), \hat{\theta} - \theta^{t+1} \rangle - \tau_\ell(\mathcal{L}_n)\mathcal{R}^2(\theta^t - \hat{\theta}).\end{aligned}\tag{77}$$

Here the second inequality follows by adding and subtracting terms.

Now for compactness in notation, define $\varphi_t(\theta) := \mathcal{L}_n(\theta^t) + \langle \nabla \mathcal{L}_n(\theta^t), \theta - \theta^t \rangle + \frac{\gamma_u}{2}\|\theta - \theta^t\|^2$, and note that by definition of the algorithm, the iterate θ^{t+1} minimizes $\varphi_t(\theta)$ over the ball $\mathbb{B}_\mathcal{R}(\rho)$. Moreover, since $\hat{\theta}$ is feasible, the first-order conditions for optimality imply that $\langle \nabla \varphi_t(\theta^{t+1}), \hat{\theta} - \theta^{t+1} \rangle \geq 0$, or equivalently that $\langle \nabla \mathcal{L}_n(\theta^t) + \gamma_u(\theta^{t+1} - \theta^t), \hat{\theta} - \theta^{t+1} \rangle \geq 0$. Applying this inequality to the lower bound (77), we find that

$$\begin{aligned}\mathcal{L}_n(\hat{\theta}) - \frac{\gamma_\ell}{2}\|\hat{\theta} - \theta^t\|^2 &\geq \mathcal{L}_n(\theta^t) + \langle \nabla \mathcal{L}_n(\theta^t), \theta^{t+1} - \theta^t \rangle + \gamma_u\langle \theta^t - \theta^{t+1}, \hat{\theta} - \theta^{t+1} \rangle - \tau_\ell(\mathcal{L}_n)\mathcal{R}^2(\theta^t - \hat{\theta}) \\ &= \varphi_t(\theta^{t+1}) - \frac{\gamma_u}{2}\|\theta^{t+1} - \theta^t\|^2 + \gamma_u\langle \theta^t - \theta^{t+1}, \hat{\theta} - \theta^{t+1} \rangle - \tau_\ell(\mathcal{L}_n)\mathcal{R}^2(\theta^t - \hat{\theta}) \\ &= \varphi_t(\theta^{t+1}) + \frac{\gamma_u}{2}\|\theta^{t+1} - \theta^t\|^2 + \gamma_u\langle \theta^t - \theta^{t+1}, \hat{\theta} - \theta^t \rangle - \tau_\ell(\mathcal{L}_n)\mathcal{R}^2(\theta^t - \hat{\theta}),\end{aligned}\tag{78}$$

where the last step follows from adding and subtracting θ^{t+1} in the inner product.

Now by the RSM condition, we have

$$\varphi_t(\theta^{t+1}) \geq \mathcal{L}_n(\theta^{t+1}) - \tau_u(\mathcal{L}_n)\mathcal{R}^2(\theta^{t+1} - \theta^t) \stackrel{(a)}{\geq} \mathcal{L}_n(\hat{\theta}) - \tau_u(\mathcal{L}_n)\mathcal{R}^2(\theta^{t+1} - \theta^t),\tag{79}$$

where inequality (a) follows by the optimality of $\hat{\theta}$, and feasibility of θ^{t+1} . Combining this inequality with the previous bound (78) yields that $\mathcal{L}_n(\hat{\theta}) - \frac{\gamma_\ell}{2}\|\hat{\theta} - \theta^t\|^2$ is lower bounded by

$$\mathcal{L}_n(\hat{\theta}) - \frac{\gamma_u}{2}\|\theta^{t+1} - \theta^t\|^2 + \gamma_u\langle \theta^t - \theta^{t+1}, \hat{\theta} - \theta^t \rangle - \tau_\ell(\mathcal{L}_n)\mathcal{R}^2(\theta^t - \hat{\theta}) - \tau_u(\mathcal{L}_n)\mathcal{R}^2(\theta^{t+1} - \theta^t),$$

and the claim (52) follows after some simple algebraic manipulations.

B Auxiliary results for Theorem 2

In this appendix, we prove the two auxiliary lemmas required in the proof of Theorem 2.

B.1 Proof of Lemma 3

This result is a generalization of an analogous result in Negahban et al. [28], with some changes required so as to adapt the statement to the optimization setting. Let θ be any vector, feasible for the problem (2), that satisfies the bound

$$\phi(\theta) \leq \phi(\theta^*) + \bar{\eta},\tag{80}$$

and assume that $\lambda_n \geq 2\mathcal{R}^*(\nabla \mathcal{L}_n(\theta^*))$. We then claim that the error vector $\Delta := \theta - \theta^*$ satisfies the inequality

$$\mathcal{R}(\Pi_{\bar{\mathcal{M}}^\perp}(\Delta)) \leq 3\mathcal{R}(\Pi_{\bar{\mathcal{M}}}(\Delta)) + 4\mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)) + 2 \min\left\{\frac{\bar{\eta}}{\lambda_n}, \bar{\rho}\right\}. \quad (81)$$

For the moment, we take this claim as given, returning later to verify its validity.

By applying this intermediate claim (81) in two different ways, we can complete the proof of Lemma 3. First, we observe that when $\theta = \hat{\theta}$, the optimality of $\hat{\theta}$ and feasibility of θ^* imply that assumption (80) holds with $\bar{\eta} = 0$, and hence the intermediate claim (81) implies that the statistical error $\Delta^* = \theta^* - \hat{\theta}$ satisfies the bound

$$\mathcal{R}(\Pi_{\bar{\mathcal{M}}^\perp}(\Delta^*)) \leq 3\mathcal{R}(\Pi_{\bar{\mathcal{M}}}(\Delta^*)) + 4\mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)). \quad (82)$$

Since $\Delta^* = \Pi_{\bar{\mathcal{M}}}(\Delta^*) + \Pi_{\bar{\mathcal{M}}^\perp}(\Delta^*)$, we can write

$$\mathcal{R}(\Delta^*) = \mathcal{R}(\Pi_{\bar{\mathcal{M}}}(\Delta^*) + \Pi_{\bar{\mathcal{M}}^\perp}(\Delta^*)) \leq 4\mathcal{R}(\Pi_{\bar{\mathcal{M}}}(\Delta^*)) + 4\mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)), \quad (83)$$

using the triangle inequality in conjunction with our earlier bound (82). Similarly, when $\theta = \theta^t$ for some $t \geq T$, then the given assumptions imply that condition (80) holds with $\bar{\eta} > 0$, so that the intermediate claim (followed by the same argument with triangle inequality) implies that the error $\Delta^t = \theta^t - \theta^*$ satisfies the bound

$$\mathcal{R}(\Delta^t) \leq 4\mathcal{R}(\Pi_{\bar{\mathcal{M}}}(\Delta^t)) + 4\mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)) + 2 \min\left\{\frac{\bar{\eta}}{\lambda_n}, \bar{\rho}\right\}. \quad (84)$$

Now let $\hat{\Delta}^t = \theta^t - \hat{\theta}$ be the optimization error at time t , and observe that we have the decomposition $\hat{\Delta}^t = \Delta^t + \Delta^*$. Consequently, by triangle inequality

$$\begin{aligned} \mathcal{R}(\hat{\Delta}^t) &\leq \mathcal{R}(\Delta^t) + \mathcal{R}(\Delta^*) \\ &\stackrel{(i)}{\leq} 4\left\{\mathcal{R}(\Pi_{\bar{\mathcal{M}}}(\Delta^t)) + \mathcal{R}(\Pi_{\bar{\mathcal{M}}}(\Delta^*))\right\} + 8\mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)) + 2 \min\left\{\frac{\bar{\eta}}{\lambda_n}, \bar{\rho}\right\} \\ &\stackrel{(ii)}{\leq} 4\Psi(\bar{\mathcal{M}}) \left\{\|\Pi_{\bar{\mathcal{M}}}(\Delta^t)\| + \|\Pi_{\bar{\mathcal{M}}}(\Delta^*)\|\right\} + 8\mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)) + 2 \min\left\{\frac{\bar{\eta}}{\lambda_n}, \bar{\rho}\right\} \\ &\stackrel{(iii)}{\leq} 4\Psi(\bar{\mathcal{M}}) \left\{\|\Delta^t\| + \|\Delta^*\|\right\} + 8\mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)) + 2 \min\left\{\frac{\bar{\eta}}{\lambda_n}, \bar{\rho}\right\}, \end{aligned} \quad (85)$$

where step (i) follows by applying both equation (83) and (84); step (ii) follows from the definition (12) of the subspace compatibility that relates the regularizer to the norm $\|\cdot\|$; and step (iii) follows from the fact that projection onto a subspace is non-expansive. Finally, since $\Delta^t = \hat{\Delta}^t - \Delta^*$, the triangle inequality implies that $\|\Delta^t\| \leq \|\hat{\Delta}^t\| + \|\Delta^*\|$. Substituting this upper bound into inequality (85) completes the proof of Lemma 3.

It remains to prove the intermediate claim (81). Letting θ be any vector, feasible for the program (2), and satisfying the condition (80), and let $\Delta = \theta - \theta^*$ be the associated error vector. Re-writing the condition (80), we have

$$\mathcal{L}_n(\theta^* + \Delta) + \lambda_n \mathcal{R}(\theta^* + \Delta) \leq \mathcal{L}_n(\theta^*) + \lambda_n \mathcal{R}(\theta^*) + \bar{\eta}.$$

Subtracting $\langle \nabla \mathcal{L}_n(\theta^*), \Delta \rangle$ from each side and then re-arranging yields the inequality

$$\mathcal{L}_n(\theta^* + \Delta) - \mathcal{L}_n(\theta^*) - \langle \nabla \mathcal{L}_n(\theta^*), \Delta \rangle + \lambda_n \{ \mathcal{R}(\theta^* + \Delta) - \mathcal{R}(\theta^*) \} \leq -\langle \nabla \mathcal{L}_n(\theta^*), \Delta \rangle + \bar{\eta}.$$

The convexity of \mathcal{L}_n then implies that $\mathcal{L}_n(\theta^* + \Delta) - \mathcal{L}_n(\theta^*) - \langle \nabla \mathcal{L}_n(\theta^*), \Delta \rangle \geq 0$, and hence that

$$\lambda_n \{ \mathcal{R}(\theta^* + \Delta) - \mathcal{R}(\theta^*) \} \leq -\langle \nabla \mathcal{L}_n(\theta^*), \Delta \rangle + \bar{\eta}.$$

Applying Hölder's inequality to $\langle \nabla \mathcal{L}_n(\theta^*), \Delta \rangle$, as expressed in terms of the dual norms \mathcal{R} and \mathcal{R}^* , yields the upper bound

$$\lambda_n \{ \mathcal{R}(\theta^* + \Delta) - \mathcal{R}(\theta^*) \} \leq \mathcal{R}^*(\nabla \mathcal{L}_n(\theta^*)) \mathcal{R}(\Delta) + \bar{\eta} \stackrel{(i)}{\leq} \frac{\lambda_n}{2} \mathcal{R}(\Delta) + \bar{\eta},$$

where step (i) uses the fact that $\lambda_n \geq 2\mathcal{R}^*(\nabla \mathcal{L}_n(\theta^*))$ by assumption.

For the remainder of the proof, let us introduce the convenient shorthand $\Delta_{\bar{\mathcal{M}}} := \Pi_{\bar{\mathcal{M}}}(\Delta)$ and $\Delta_{\bar{\mathcal{M}}^\perp} := \Pi_{\bar{\mathcal{M}}^\perp}(\Delta)$, with similar shorthand for projections involving θ^* . Making note of the decomposition $\Delta = \Delta_{\bar{\mathcal{M}}} + \Delta_{\bar{\mathcal{M}}^\perp}$, an application of triangle inequality then yields the upper bound

$$\mathcal{R}(\theta^* + \Delta) - \mathcal{R}(\theta^*) \leq \frac{1}{2} \{ \mathcal{R}(\Delta_{\bar{\mathcal{M}}}) + \mathcal{R}(\Delta_{\bar{\mathcal{M}}^\perp}) \} + \frac{\bar{\eta}}{\lambda_n}, \quad (86)$$

where we have rescaled both sides by $\lambda_n > 0$.

It remains to further lower bound the left-hand side (86). By triangle inequality, we have

$$-\mathcal{R}(\theta^*) \geq -\mathcal{R}(\theta_{\mathcal{M}}^*) - \mathcal{R}(\theta_{\mathcal{M}}^*). \quad (87)$$

Let us now write $\theta^* + \Delta = \theta_{\mathcal{M}}^* + \theta_{\mathcal{M}}^* + \Delta_{\bar{\mathcal{M}}} + \Delta_{\bar{\mathcal{M}}^\perp}$. Using this representation and triangle inequality, we have

$$\mathcal{R}(\theta^* + \Delta) \geq \mathcal{R}(\theta_{\mathcal{M}}^* + \Delta_{\bar{\mathcal{M}}^\perp}) - \mathcal{R}(\theta_{\mathcal{M}}^* + \Delta_{\bar{\mathcal{M}}}) \geq \mathcal{R}(\theta_{\mathcal{M}}^* + \Delta_{\bar{\mathcal{M}}^\perp}) - \mathcal{R}(\theta_{\mathcal{M}}^*) - \mathcal{R}(\Delta_{\bar{\mathcal{M}}}).$$

Finally, since $\theta_{\mathcal{M}}^* \in \mathcal{M}$ and $\Delta_{\bar{\mathcal{M}}^\perp} \in \bar{\mathcal{M}}^\perp$, the decomposability of \mathcal{R} implies that $\mathcal{R}(\theta_{\mathcal{M}}^* + \Delta_{\bar{\mathcal{M}}^\perp}) = \mathcal{R}(\theta_{\mathcal{M}}^*) + \mathcal{R}(\Delta_{\bar{\mathcal{M}}^\perp})$, and hence that

$$\mathcal{R}(\theta^* + \Delta) \geq \mathcal{R}(\theta_{\mathcal{M}}^*) + \mathcal{R}(\Delta_{\bar{\mathcal{M}}^\perp}) - \mathcal{R}(\theta_{\mathcal{M}}^*) - \mathcal{R}(\Delta_{\bar{\mathcal{M}}}). \quad (88)$$

Adding together equations (87) and (88), we obtain the lower bound

$$\mathcal{R}(\theta^* + \Delta) - \mathcal{R}(\theta^*) \geq \mathcal{R}(\Delta_{\bar{\mathcal{M}}^\perp}) - 2\mathcal{R}(\theta_{\mathcal{M}}^*) - \mathcal{R}(\Delta_{\bar{\mathcal{M}}}). \quad (89)$$

Combining this lower bound with the earlier inequality (86), some algebra yields the bound

$$\mathcal{R}(\Delta_{\bar{\mathcal{M}}^\perp}) \leq 3\mathcal{R}(\Delta_{\bar{\mathcal{M}}}) + 4\mathcal{R}(\theta_{\mathcal{M}}^*) + 2\frac{\eta}{\lambda_n},$$

corresponding to the bound (81) when η/λ_n achieves the final minimum. To obtain the final term involving $\bar{\rho}$ in the bound (81), two applications of triangle inequality yields

$$\mathcal{R}(\Delta_{\bar{\mathcal{M}}^\perp}) \leq \mathcal{R}(\Delta_{\bar{\mathcal{M}}}) + \mathcal{R}(\Delta) \leq \mathcal{R}(\Delta_{\bar{\mathcal{M}}}) + 2\bar{\rho},$$

where we have used the fact that $\mathcal{R}(\Delta) \leq \mathcal{R}(\theta) + \mathcal{R}(\theta^*) \leq 2\bar{\rho}$, since both θ and θ^* are feasible for the program (2).

B.2 Proof of Lemma 4

The proof of this result follows lines similar to the proof of convergence by Nesterov [32]. Recall our notation $\phi(\theta) = \mathcal{L}_n(\theta) + \lambda_n \mathcal{R}(\theta)$, $\widehat{\Delta}^t = \theta^t - \widehat{\theta}$, and that $\eta_\phi^t = \phi(\theta^t) - \phi(\widehat{\theta})$. We begin by proving that under the stated conditions, a useful version of restricted strong convexity (48) is in force:

Lemma 11. *Under the assumptions of Lemma 4, we are guaranteed that*

$$\left\{ \frac{\gamma_\ell}{2} - 32\tau_\ell(\mathcal{L}_n)\Psi^2(\overline{\mathcal{M}}) \right\} \|\widehat{\Delta}^t\|^2 \leq 2\tau_\ell(\mathcal{L}_n)v^2 + \phi(\theta^t) - \phi(\widehat{\theta}), \quad \text{and} \quad (90a)$$

$$\left\{ \frac{\gamma_\ell}{2} - 32\tau_\ell(\mathcal{L}_n)\Psi^2(\overline{\mathcal{M}}) \right\} \|\widehat{\Delta}^t\|^2 \leq 2\tau_\ell(\mathcal{L}_n)v^2 + \mathcal{T}_\mathcal{L}(\widehat{\theta}; \theta^t), \quad (90b)$$

where $v := \bar{\epsilon}_{\text{stat}} + 2 \min\left(\frac{\bar{\eta}}{\lambda_n}, \bar{\rho}\right)$.

See Appendix B.3 for the proof of this claim. So as to ease notation in the remainder of the proof, let us introduce the shorthand

$$\phi_t(\theta) := \mathcal{L}_n(\theta^t) + \langle \nabla \mathcal{L}_n(\theta^t), \theta - \theta^t \rangle + \frac{\gamma_u}{2} \|\theta - \theta^t\|^2 + \lambda_n \mathcal{R}(\theta), \quad (91)$$

corresponding to the approximation to the regularized loss function ϕ that is minimized at iteration t of the update (4). Since θ^{t+1} minimizes ϕ_t over the set $\mathbb{B}_\mathcal{R}(\bar{\rho})$, we are guaranteed that $\phi_t(\theta^{t+1}) \leq \phi_t(\theta)$ for all $\theta \in \mathbb{B}_\mathcal{R}(\bar{\rho})$. In particular, for any $\alpha \in (0, 1)$, the vector $\theta_\alpha = \alpha\widehat{\theta} + (1 - \alpha)\theta^t$ lies in the convex set $\mathbb{B}_\mathcal{R}(\bar{\rho})$, so that

$$\begin{aligned} \phi_t(\theta^{t+1}) &\leq \phi_t(\theta_\alpha) = \mathcal{L}_n(\theta^t) + \langle \nabla \mathcal{L}_n(\theta^t), \theta_\alpha - \theta^t \rangle + \frac{\gamma_u}{2} \|\theta_\alpha - \theta^t\|^2 + \lambda_n \mathcal{R}(\theta_\alpha) \\ &\stackrel{(i)}{=} \mathcal{L}_n(\theta^t) + \langle \nabla \mathcal{L}_n(\theta^t), \alpha\widehat{\theta} - \alpha\theta^t \rangle + \frac{\gamma_u \alpha^2}{2} \|\widehat{\theta} - \theta^t\|^2 + \lambda_n \mathcal{R}(\theta_\alpha) \\ &\stackrel{(ii)}{\leq} \mathcal{L}_n(\theta^t) + \langle \nabla \mathcal{L}_n(\theta^t), \alpha\widehat{\theta} - \alpha\theta^t \rangle + \frac{\gamma_u \alpha^2}{2} \|\widehat{\theta} - \theta^t\|^2 + \lambda_n \alpha \mathcal{R}(\widehat{\theta}) + \lambda_n (1 - \alpha) \mathcal{R}(\theta^t), \end{aligned}$$

where step (i) follows from substituting the definition of θ_α , and step (ii) uses the convexity of the regularizer \mathcal{R} .

Now, the stated conditions of the lemma ensure that $\gamma_\ell/2 - 32\tau_\ell(\mathcal{L}_n)\Psi^2(\overline{\mathcal{M}}) \geq 0$, so that by equation (90b), we have $\mathcal{L}_n(\widehat{\theta}) + 2\tau_\ell(\mathcal{L}_n)v^2 \geq \mathcal{L}_n(\theta^t) + \langle \nabla \mathcal{L}_n(\theta^t), \widehat{\theta} - \theta^t \rangle$. Substituting back into our earlier bound yields

$$\begin{aligned} \phi_t(\theta^{t+1}) &\leq (1 - \alpha)\mathcal{L}_n(\theta^t) + \alpha\mathcal{L}_n(\widehat{\theta}) + 2\alpha\tau_\ell(\mathcal{L}_n)v^2 + \frac{\gamma_u \alpha^2}{2} \|\widehat{\theta} - \theta^t\|^2 + \alpha\lambda_n \mathcal{R}(\widehat{\theta}) + (1 - \alpha)\lambda_n \mathcal{R}(\theta^t) \\ &\stackrel{(iii)}{=} \phi(\theta^t) - \alpha(\phi(\theta^t) - \phi(\widehat{\theta})) + 2\tau_\ell(\mathcal{L}_n)v^2 + \frac{\gamma_u \alpha^2}{2} \|\widehat{\theta} - \theta^t\|^2, \end{aligned} \quad (92)$$

where we have used the definition of ϕ and $\alpha \leq 1$ in step (iii).

In order to complete the proof, it remains to relate $\phi_t(\theta^{t+1})$ to $\phi(\theta^{t+1})$, which can be performed by exploiting restricted smoothness. In particular, applying the RSM condition at the iterate θ^{t+1} in the direction θ^t yields the upper bound

$$\mathcal{L}_n(\theta^{t+1}) \leq \mathcal{L}_n(\theta^t) + \langle \mathcal{L}_n(\theta^t), \theta^{t+1} - \theta^t \rangle + \frac{\gamma_u}{2} \|\theta^{t+1} - \theta^t\|^2 + \tau_u(\mathcal{L}_n) \mathcal{R}^2(\theta^{t+1} - \theta^t),$$

so that

$$\begin{aligned}\phi(\theta^{t+1}) &\leq \mathcal{L}_n(\theta^t) + \langle \mathcal{L}_n(\theta^t), \theta^{t+1} - \theta^t \rangle + \frac{\gamma_u}{2} \|\theta^{t+1} - \theta^t\|^2 + \tau_u(\mathcal{L}_n) \mathcal{R}^2(\theta^{t+1} - \theta^t) + \lambda_n \mathcal{R}(\theta^{t+1}) \\ &= \phi_t(\theta^{t+1}) + \tau_u(\mathcal{L}_n) \mathcal{R}^2(\theta^{t+1} - \theta^t).\end{aligned}$$

Combining the above bound with the inequality (92) and recalling the notation $\widehat{\Delta}^t = \theta^t - \widehat{\theta}$, we obtain

$$\begin{aligned}\phi(\theta^{t+1}) &\leq \phi(\theta^t) - \alpha(\phi(\theta^t) - \phi(\widehat{\theta})) + \frac{\gamma_u \alpha^2}{2} \|\widehat{\theta} - \theta^t\|^2 + \tau_u(\mathcal{L}_n) \mathcal{R}^2(\theta^{t+1} - \theta^t) + 2\tau_\ell(\mathcal{L}_n) v^2 \\ &\stackrel{(iv)}{\leq} \phi(\theta^t) - \alpha(\phi(\theta^t) - \phi(\widehat{\theta})) + \frac{\gamma_u \alpha^2}{2} \|\widehat{\Delta}^t\|^2 + \tau_u(\mathcal{L}_n) [\mathcal{R}(\widehat{\Delta}^{t+1}) + \mathcal{R}(\widehat{\Delta}^t)]^2 + 2\tau_\ell(\mathcal{L}_n) v^2 \\ &\stackrel{(v)}{\leq} \phi(\theta^t) - \alpha(\phi(\theta^t) - \phi(\widehat{\theta})) + \frac{\gamma_u \alpha^2}{2} \|\widehat{\Delta}^t\|^2 + 2\tau_u(\mathcal{L}_n) (\mathcal{R}^2(\widehat{\Delta}^{t+1}) + \mathcal{R}^2(\widehat{\Delta}^t)) + 2\tau_\ell(\mathcal{L}_n) v^2.\end{aligned}\tag{93}$$

Here step (iv) uses the fact that $\theta^t - \theta^{t+1} = \widehat{\Delta}^t - \widehat{\Delta}^{t+1}$ and applies triangle inequality to the norm \mathcal{R} , whereas step (v) follows from Cauchy-Schwarz inequality.

Next, combining Lemma 3 with the Cauchy-Schwarz inequality yields the upper bound

$$\mathcal{R}^2(\widehat{\Delta}^t) \leq 32\Psi^2(\overline{\mathcal{M}}) \|\widehat{\Delta}^t\|^2 + 2v^2\tag{94}$$

where $v = \bar{\epsilon}_{\text{stat}}(\mathcal{M}, \overline{\mathcal{M}}) + 2 \min(\frac{\bar{\eta}}{\lambda_n}, \bar{\rho})$, is a constant independent of θ^t and $\bar{\epsilon}_{\text{stat}}(\mathcal{M}, \overline{\mathcal{M}})$ was previously defined in the lemma statement. Substituting the above bound into inequality (93) yields that $\phi(\theta^{t+1})$ is at most

$$\begin{aligned}\phi(\theta^t) - \alpha(\phi(\theta^t) - \phi(\widehat{\theta})) + \frac{\gamma_u \alpha^2}{2} \|\widehat{\Delta}^t\|^2 &+ 64\tau_u(\mathcal{L}_n) \Psi^2(\overline{\mathcal{M}}) \|\widehat{\Delta}^{t+1}\|^2 \\ &+ 64\tau_u(\mathcal{L}_n) \Psi^2(\overline{\mathcal{M}}) \|\widehat{\Delta}^t\|^2 + 8\tau_u(\mathcal{L}_n) v^2 + 2\tau_\ell(\mathcal{L}_n) v^2.\end{aligned}\tag{95}$$

The final step is to translate quantities involving $\widehat{\Delta}^t$ to functional values, which may be done using the RSC condition (90a) from Lemma 11. In particular, combining the RSC condition (90a) with the inequality (95) yields

$$\begin{aligned}\phi(\theta^{t+1}) &\leq \phi(\theta^t) - \alpha \eta_\phi^t + \frac{(\gamma_u \alpha^2 + 64\tau_u(\mathcal{L}_n) \Psi^2(\overline{\mathcal{M}}))}{\bar{\gamma}_\ell} (\eta_\phi^t + 2\tau_\ell(\mathcal{L}_n) v^2) + \\ &\quad \frac{64\tau_u(\mathcal{L}_n) \Psi^2(\overline{\mathcal{M}})}{\bar{\gamma}_\ell} (\eta_\phi^{t+1} + 2\tau_\ell(\mathcal{L}_n) v^2) + 8\tau_u(\mathcal{L}_n) v^2 + 2\tau_\ell(\mathcal{L}_n) v^2.\end{aligned}$$

where we have introduced the shorthand $\bar{\gamma}_\ell := \gamma_\ell - 64\tau_\ell(\mathcal{L}_n) \Psi^2(\overline{\mathcal{M}})$. Recalling the definition of β , adding and subtracting $\phi(\widehat{\theta})$ from both sides, and choosing $\alpha = \frac{\bar{\gamma}_\ell}{2\gamma_u} \in (0, 1)$, we obtain

$$\left(1 - \frac{64\tau_u(\mathcal{L}_n) \Psi^2(\overline{\mathcal{M}})}{\bar{\gamma}_\ell}\right) \eta_\phi^{t+1} \leq \left(1 - \frac{\bar{\gamma}_\ell}{4\gamma_u} + \frac{64\tau_u(\mathcal{L}_n) \Psi^2(\overline{\mathcal{M}})}{\bar{\gamma}_\ell}\right) \eta_\phi^t + \beta(\overline{\mathcal{M}}) v^2.$$

Recalling the definition of the contraction factor κ from the statement of Theorem 2, the above expression can be rewritten as

$$\eta_\phi^{t+1} \leq \kappa \eta_\phi^t + \beta(\overline{\mathcal{M}}) \xi(\overline{\mathcal{M}}) v^2, \quad \text{where } \xi(\mathcal{M}) = \left\{1 - \frac{64\tau_u(\mathcal{L}_n) \Psi^2(\overline{\mathcal{M}})}{\bar{\gamma}_\ell}\right\}^{-1}.$$

Finally, iterating the above expression yields $\eta_\phi^t \leq \kappa^{t-T} \eta_\phi^T + \frac{\xi(\overline{\mathcal{M}})\beta(\overline{\mathcal{M}})v^2}{1-\kappa}$, where we have used the condition $\kappa \in (0, 1)$ in order to sum the geometric series, thereby completing the proof.

B.3 Proof of Lemma 11

The key idea to prove the lemma is to use the definition of RSC along with the iterated cone bound of Lemma 3 for simplifying the error terms in RSC.

Let us first show that condition (90a) holds. From the RSC condition assumed in the lemma statement, we have

$$\mathcal{L}_n(\theta^t) - \mathcal{L}_n(\widehat{\theta}) - \langle \nabla \mathcal{L}_n(\widehat{\theta}), \theta^t - \widehat{\theta} \rangle \geq \frac{\gamma_\ell}{2} \|\widehat{\theta} - \theta^t\|^2 - \tau_\ell(\mathcal{L}_n) \mathcal{R}^2(\widehat{\theta} - \theta^t). \quad (96)$$

From the convexity of \mathcal{R} and definition of the subdifferential $\partial\mathcal{R}(\theta)$, we obtain

$$\mathcal{R}(\theta^t) - \mathcal{R}(\widehat{\theta}) - \langle \partial\mathcal{R}(\widehat{\theta}), \theta^t - \widehat{\theta} \rangle \geq 0.$$

Adding this lower bound with the inequality (96) yields

$$\phi(\theta^t) - \phi(\widehat{\theta}) - \langle \nabla \phi(\widehat{\theta}), \theta^t - \widehat{\theta} \rangle \geq \frac{\gamma_\ell}{2} \|\widehat{\theta} - \theta^t\|^2 - \tau_\ell(\mathcal{L}_n) \mathcal{R}^2(\widehat{\theta} - \theta^t),$$

where we recall that $\phi(\theta) = \mathcal{L}_n(\theta) + \lambda_n \mathcal{R}(\theta)$ is our objective function. By the optimality of $\widehat{\theta}$ and feasibility of θ^t , we are guaranteed that $\langle \nabla \phi(\widehat{\theta}), \theta^t - \widehat{\theta} \rangle \geq 0$, and hence

$$\begin{aligned} \phi(\theta^t) - \phi(\widehat{\theta}) &\geq \frac{\gamma_\ell}{2} \|\widehat{\theta} - \theta^t\|^2 - \tau_\ell(\mathcal{L}_n) \mathcal{R}^2(\widehat{\theta} - \theta^t) \\ &\stackrel{(i)}{\geq} \frac{\gamma_\ell}{2} \|\widehat{\theta} - \theta^t\|^2 - \tau_\ell(\mathcal{L}_n) \{32\Psi^2(\overline{\mathcal{M}}) \|\widehat{\theta} - \theta^t\|^2 + 2v^2\} \end{aligned}$$

where step (i) follows by applying Lemma 3. Some algebra then yields the claim (90a).

Finally, let us verify the claim (90b). Using the RSC condition, we have

$$\mathcal{L}_n(\widehat{\theta}) - \mathcal{L}_n(\theta^t) - \langle \nabla \mathcal{L}_n(\theta^t), \widehat{\theta} - \theta^t \rangle \geq \frac{\gamma_\ell}{2} \|\widehat{\theta} - \theta^t\|^2 - \tau_\ell(\mathcal{L}_n) \mathcal{R}^2(\widehat{\theta} - \theta^t). \quad (97)$$

As before, applying Lemma 3 yields

$$\underbrace{\mathcal{L}_n(\widehat{\theta}) - \mathcal{L}_n(\theta^t) - \langle \nabla \mathcal{L}_n(\theta^t), \widehat{\theta} - \theta^t \rangle}_{\tau_{\mathcal{L}}(\widehat{\theta}; \theta^t)} \geq \frac{\gamma_\ell}{2} \|\widehat{\theta} - \theta^t\|^2 - \tau_\ell(\mathcal{L}_n) \left(32\Psi^2(\overline{\mathcal{M}}) \|\widehat{\theta} - \theta^t\|^2 + 2v^2\right),$$

and rearranging the terms and establishes the claim (90b).

C Proof of Lemma 5

Given the condition $\mathcal{R}(\widehat{\theta}) \leq \rho \leq \mathcal{R}(\theta^*)$, we have $\mathcal{R}(\widehat{\theta}) = \mathcal{R}(\theta^* + \Delta^*) \leq \mathcal{R}(\theta^*)$. By triangle inequality, we have

$$\mathcal{R}(\theta^*) = \mathcal{R}(\Pi_{\mathcal{M}}(\theta^*) + \Pi_{\mathcal{M}^\perp}(\theta^*)) \leq \mathcal{R}(\Pi_{\mathcal{M}}(\theta^*)) + \mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)).$$

We then write

$$\begin{aligned}
\mathcal{R}(\theta^* + \Delta^*) &= \mathcal{R}(\Pi_{\mathcal{M}}(\theta^*) + \Pi_{\mathcal{M}^\perp}(\theta^*) + \Pi_{\bar{\mathcal{M}}}(\Delta^*) + \Pi_{\bar{\mathcal{M}}^\perp}(\Delta^*)) \\
&\stackrel{(i)}{\geq} \mathcal{R}(\Pi_{\mathcal{M}}(\theta^*) + \Pi_{\bar{\mathcal{M}}^\perp}(\Delta^*)) - \mathcal{R}(\Pi_{\bar{\mathcal{M}}}(\Delta^*)) - \mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)) \\
&\stackrel{(ii)}{=} \mathcal{R}(\Pi_{\mathcal{M}}(\theta^*)) + \mathcal{R}(\Pi_{\bar{\mathcal{M}}^\perp}(\Delta^*)) - \mathcal{R}(\Pi_{\bar{\mathcal{M}}}(\Delta^*)) - \mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)),
\end{aligned}$$

where the bound (i) follows by triangle inequality, and step (ii) uses the decomposability of \mathcal{R} over the pair \mathcal{M} and $\bar{\mathcal{M}}^\perp$. By combining this lower bound with the previously established upper bound

$$\mathcal{R}(\theta^* + \Delta^*) \leq \mathcal{R}(\Pi_{\mathcal{M}}(\theta^*)) + \mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)),$$

we conclude that $\mathcal{R}(\Pi_{\bar{\mathcal{M}}^\perp}(\Delta^*)) \leq \mathcal{R}(\Pi_{\bar{\mathcal{M}}}(\Delta^*)) + 2\mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*))$. Finally, by triangle inequality, we have $\mathcal{R}(\Delta^*) \leq \mathcal{R}(\Pi_{\bar{\mathcal{M}}}(\Delta^*)) + \mathcal{R}(\Pi_{\bar{\mathcal{M}}^\perp}(\Delta^*))$, and hence

$$\begin{aligned}
\mathcal{R}(\Delta^*) &\leq 2\mathcal{R}(\Pi_{\bar{\mathcal{M}}}(\Delta^*)) + 2\mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)) \\
&\stackrel{(i)}{\leq} 2\Psi(\bar{\mathcal{M}}^\perp)\|\Pi_{\bar{\mathcal{M}}}(\Delta^*)\| + 2\mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)) \\
&\stackrel{(ii)}{\leq} 2\Psi(\bar{\mathcal{M}}^\perp)\|\Delta^*\| + 2\mathcal{R}(\Pi_{\mathcal{M}^\perp}(\theta^*)),
\end{aligned}$$

where inequality (i) follows from Definition 4 of the subspace compatibility Ψ , and the bound (ii) follows from non-expansivity of projection onto a subspace.

D A general result on Gaussian observation operators

In this appendix, we state a general result about a Gaussian random matrices, and show how it can be adapted to prove Lemmas 6 and 7. Let $X \in \mathbb{R}^{n \times d}$ be a Gaussian random matrix with i.i.d. rows $x_i \sim N(0, \Sigma)$, where $\Sigma \in \mathbb{R}^{d \times d}$ is a covariance matrix. We refer to X as a sample from the Σ -Gaussian ensemble. In order to state the result, we use $\Sigma^{1/2}$ to denote the symmetric matrix square root.

Proposition 1. *Given a random matrix X drawn from the Σ -Gaussian ensemble, there are universal constants c_i , $i = 0, 1$ such that*

$$\frac{\|X\theta\|_2^2}{n} \geq \frac{1}{2}\|\Sigma^{1/2}\theta\|_2^2 - c_1 \frac{(\mathbb{E}[\mathcal{R}^*(x_i)])^2}{n} \mathcal{R}^2(\theta) \quad \text{and} \quad (98a)$$

$$\frac{\|X\theta\|_2^2}{n} \leq 2\|\Sigma^{1/2}\theta\|_2^2 + c_1 \frac{(\mathbb{E}[\mathcal{R}^*(x_i)])^2}{n} \mathcal{R}^2(\theta) \quad \text{for all } \theta \in \mathbb{R}^d \quad (98b)$$

with probability greater than $1 - \exp(-c_0 n)$.

We omit the proof of this result. The two special instances proved in Lemma 6 and 7 have been proved in the papers [35] and [29] respectively. We now show how Proposition 1 can be used to recover various lemmas required in our proofs.

Proof of Lemma 6: We begin by establishing this auxiliary result required in the proof of Corollary 2. When $\mathcal{R}(\cdot) = \|\cdot\|_1$, we have $\mathcal{R}^*(\cdot) = \|\cdot\|_\infty$. Moreover, the random vector $x_i \sim N(0, \Sigma)$ can be written as $x_i = \Sigma^{1/2}w$, where $w \sim N(0, I_{d \times d})$ is standard normal. Consequently, using properties of Gaussian maxima [23] and defining $\zeta(\Sigma) = \max_{j=1,2,\dots,d} \Sigma_{jj}$, we have the bound

$$(\mathbb{E}[\|x_i\|_\infty])^2 \leq \zeta(\Sigma) (\mathbb{E}[\|w\|_\infty])^2 \leq 3\zeta(\Sigma) \sqrt{\log d}.$$

Substituting into Proposition 1 yields the claims (62a) and (62b).

Proof of Lemma 7: In order to prove this claim, we view each random observation matrix $X_i \in \mathbb{R}^{d \times d}$ as a $d = d^2$ vector (namely the quantity $\text{vec}(X_i)$), and apply Proposition 1 in this vectorized setting. Given the standard Gaussian vector $w \in \mathbb{R}^{d^2}$, we let $W \in \mathbb{R}^{d \times d}$ be the random matrix such that $\text{vec}(W) = w$. With this notation, the term $\mathcal{R}^*(\text{vec}(X_i))$ is equivalent to the operator norm $\|X_i\|_{\text{op}}$. As shown in Negahban and Wainwright [29], $\mathbb{E}[\|X_i\|_{\text{op}}] \leq 24\zeta_{\text{mat}}(\Sigma) \sqrt{d}$, where ζ_{mat} was previously defined (65).

E Auxiliary results for Corollary 5

In this section, we provide the proofs of Lemmas 8 and 9 that play a central role in the proof of Corollary 5. In order to do so, we require the following result, which is a re-statement of a theorem due to Negahban and Wainwright [30]:

Proposition 2. *For the matrix completion operator \mathfrak{X}_n , there are universal positive constants (c_1, c_2) such that*

$$\left| \frac{\|\mathfrak{X}_n(\Theta)\|_2^2}{n} - \|\Theta\|_F^2 \right| \leq c_1 d \|\Theta\|_\infty \|\Theta\|_1 \sqrt{\frac{d \log d}{n}} + c_2 \left(d \|\Theta\|_\infty \sqrt{\frac{d \log d}{n}} \right)^2 \quad \text{for all } \Theta \in \mathbb{R}^{d \times d} \quad (99)$$

with probability at least $1 - \exp(-d \log d)$.

E.1 Proof of Lemma 8

Applying Proposition 2 to $\widehat{\Delta}^t$ and using the fact that $d\|\widehat{\Delta}^t\|_\infty \leq 2\alpha$ yields

$$\frac{\|\mathfrak{X}_n(\widehat{\Delta}^t)\|_2^2}{n} \geq \|\widehat{\Delta}^t\|_F^2 - c_1 \alpha \|\widehat{\Delta}^t\|_1 \sqrt{\frac{d \log d}{n}} - c_2 \alpha^2 \frac{d \log d}{n}, \quad (100)$$

where we recall our convention of allowing the constants to change from line to line. From Lemma 1,

$$\|\widehat{\Delta}^t\|_1 \leq 2\Psi(\overline{\mathcal{M}}^\perp) \|\widehat{\Delta}^t\|_F + 2\|\Pi_{\mathcal{M}^\perp}(\theta^*)\|_1 + 2\|\Delta^*\|_1 + \Psi(\overline{\mathcal{M}}^\perp) \|\Delta^*\|_F.$$

Since $\rho \leq \|\Theta^*\|_1$, Lemma 5 implies that $\|\Delta^*\|_1 \leq 2\Psi(\overline{\mathcal{M}}^\perp) \|\Delta^*\|_F + \|\Pi_{\mathcal{M}^\perp}(\theta^*)\|_1$, and hence that

$$\|\widehat{\Delta}^t\|_1 \leq 2\Psi(\overline{\mathcal{M}}^\perp) \|\widehat{\Delta}^t\|_F + 4\|\Pi_{\mathcal{M}^\perp}(\theta^*)\|_1 + 5\Psi(\overline{\mathcal{M}}^\perp) \|\Delta^*\|_F. \quad (101)$$

Combined with the lower bound, we obtain that $\frac{\|\mathfrak{X}_n(\widehat{\Delta}^t)\|_2^2}{n}$ is lower bounded by

$$\|\widehat{\Delta}^t\|_F^2 \left\{ 1 - \frac{2c_1 \alpha \Psi(\mathcal{M}^\perp) \sqrt{\frac{d \log d}{n}}}{\|\widehat{\Delta}^t\|_F} \right\} - 2c_1 \alpha \sqrt{\frac{d \log d}{n}} \left\{ 4\|\Pi_{\mathcal{M}^\perp}(\theta^*)\|_1 + 5\Psi(\mathcal{M}^\perp)\|\Delta^*\|_F \right\} - c_2 \alpha^2 \frac{d \log d}{n}.$$

Consequently, for all iterations such that $\|\widehat{\Delta}^t\|_F \geq 4c_1 \Psi(\mathcal{M}^\perp) \sqrt{\frac{d \log d}{n}}$, we have

$$\frac{\|\mathfrak{X}_n(\widehat{\Delta}^t)\|_2^2}{n} \geq \frac{1}{2} \|\widehat{\Delta}^t\|_F^2 - 2c_1 \alpha \sqrt{\frac{d \log d}{n}} \left\{ 4\|\Pi_{\mathcal{M}^\perp}(\theta^*)\|_1 + 5\Psi(\mathcal{M}^\perp)\|\Delta^*\|_F \right\} - c_2 \alpha^2 \frac{d \log d}{n}.$$

By subtracting off an additional term, the bound is valid for all $\widehat{\Delta}^t$ —viz.

$$\begin{aligned} \frac{\|\mathfrak{X}_n(\widehat{\Delta}^t)\|_2^2}{n} &\geq \frac{1}{2} \|\widehat{\Delta}^t\|_F^2 - 2c_1 \alpha \sqrt{\frac{d \log d}{n}} \left\{ 4\|\Pi_{\mathcal{M}^\perp}(\theta^*)\|_1 + 5\Psi(\mathcal{M}^\perp)\|\Delta^*\|_F \right\} \\ &\quad - c_2 \alpha^2 \frac{d \log d}{n} - 16c_1^2 \alpha^2 \Psi^2(\mathcal{M}^\perp) \frac{d \log d}{n}. \end{aligned}$$

E.2 Proof of Lemma 9

Applying Proposition 2 to Γ^t and using the fact that $d\|\Gamma^t\|_\infty \leq 2\alpha$ yields

$$\frac{\|\mathfrak{X}_n(\Gamma^t)\|_2^2}{n} \leq \|\Gamma^t\|_F^2 + c_1 \alpha \|\Gamma^t\|_1 \sqrt{\frac{d \log d}{n}} + c_2 \alpha^2 \frac{d \log d}{n}, \quad (102)$$

where we recall our convention of allowing the constants to change from line to line. By triangle inequality, we have $\|\Gamma^t\|_1 \leq \|\Theta^t - \widehat{\Theta}\|_1 + \|\Theta^{t+1} - \widehat{\Theta}\|_1 = \|\widehat{\Delta}^t\|_1 + \|\widehat{\Delta}^{t+1}\|_1$. Equation 101 gives us bounds on $\|\widehat{\Delta}^t\|_1$ and $\|\widehat{\Delta}^{t+1}\|_1$. Substituting them into the upper bound (102) yields the claim.

References

- [1] A. Agarwal, S. Negahban, and M. J. Wainwright. Noisy matrix decomposition via convex relaxation: Optimal rates in high dimensions. *To appear in Annals of Statistics*, 2011. Appeared as <http://arxiv.org/abs/1102.4807>.
- [2] A. A. Amini and M. J. Wainwright. High-dimensional analysis of semidefinite relaxations for sparse principal component analysis. *Annals of Statistics*, 37:2877–2921, 2009.
- [3] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.
- [4] S. Becker, J. Bobin, and E. J. Candes. Nesta: a fast and accurate first-order method for sparse recovery. *SIAM Journal on Imaging Sciences*, 4(1):1–39, 2011.
- [5] D.P. Bertsekas. *Nonlinear programming*. Athena Scientific, Belmont, MA, 1995.
- [6] P. J. Bickel, Y. Ritov, and A. Tsybakov. Simultaneous analysis of Lasso and Dantzig selector. *Annals of Statistics*, 37(4):1705–1732, 2009.

- [7] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, Cambridge, UK, 2004.
- [8] K. Bredies and D. A. Lorenz. Linear convergence of iterative soft-thresholding. *Journal of Fourier Analysis and Applications*, 14:813–837, 2008.
- [9] F. Bunea, A. Tsybakov, and M. Wegkamp. Sparsity oracle inequalities for the Lasso. *Electronic Journal of Statistics*, pages 169–194, 2007.
- [10] E. J. Candes, X. Li, Y. Ma, and J. Wright. Robust Principal Component Analysis? *J. ACM*, 58:11:1–11:37, 2011.
- [11] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Found. Comput. Math.*, 9(6):717–772, 2009.
- [12] V. Chandrasekaran, S. Sanghavi, P. Parrilo, and A. Willsky. Rank-sparsity incoherence for matrix decomposition. *SIAM J. on Optimization*, 21(2):572–596, 2011.
- [13] S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by basis pursuit. *SIAM J. Sci. Computing*, 20(1):33–61, 1998.
- [14] J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra. Efficient projections onto the ℓ_1 -ball for learning in high dimensions. In *ICML*, 2008.
- [15] R. L. Dykstra. An iterative procedure for obtaining i-projections onto the intersection of convex sets. *Annals of Probability*, 13(3):975–984, 1985.
- [16] M. Fazel. *Matrix Rank Minimization with Applications*. PhD thesis, Stanford, 2002. Available online: <http://faculty.washington.edu/mfazel/thesis-final.pdf>.
- [17] R. Garg and R. Khandekar. Gradient descent with sparsification: an iterative algorithm for sparse recovery with restricted isometry property. In *ICML*, 2009.
- [18] E. T. Hale, Y. Wotao, and Y. Zhang. Fixed-point continuation for ℓ_1 -minimization: Methodology and convergence. *SIAM J. on Optimization*, 19(3):1107–1130, 2008.
- [19] D. Hsu, S.M. Kakade, and Tong Zhang. Robust matrix decomposition with sparse corruptions. *IEEE Trans. Info. Theory*, 57(11):7221 –7234, 2011.
- [20] J. Huang and T. Zhang. The benefit of group sparsity. *The Annals of Statistics*, 38(4):1978–2004, 2010.
- [21] S. Ji and J. Ye. An accelerated gradient method for trace norm minimization. In *ICML*, 2009.
- [22] V. Koltchinskii, K. Lounici, and A. B. Tsybakov. Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *Annals of Statistics*, 39:2302–2329, 2011.
- [23] M. Ledoux and M. Talagrand. *Probability in Banach Spaces: Isoperimetry and Processes*. Springer-Verlag, New York, NY, 1991.

- [24] K. Lee and Y. Bresler. Guaranteed minimum rank approximation from linear observations by nuclear norm minimization with an ellipsoidal constraint. Technical report, UIUC, 2009. Available at arXiv:0903.4742.
- [25] K. Lounici, M. Pontil, A. B. Tsybakov, and S. van de Geer. Taking advantage of sparsity in multi-task learning. In *COLT*, 2009.
- [26] Z. Q. Luo and P. Tseng. Error bounds and convergence analysis of feasible descent methods: a general approach. *Annals of Operations Research*, 46-47:157–178, 1993.
- [27] N. Meinshausen and P. Bühlmann. High-dimensional graphs and variable selection with the Lasso. *Annals of Statistics*, 34:1436–1462, 2006.
- [28] S. Negahban, P. Ravikumar, M. J. Wainwright, and B. Yu. A unified framework for high-dimensional analysis of M-estimators with decomposable regularizers. In *NIPS*, 2009. To appear in *Statistical Science*.
- [29] S. Negahban and M. J. Wainwright. Estimation of (near) low-rank matrices with noise and high-dimensional scaling. *Annals of Statistics*, 39(2):1069–1097, 2011.
- [30] S. Negahban and M. J. Wainwright. Restricted strong convexity and (weighted) matrix completion: Optimal bounds with noise. *Journal of Machine Learning Research*, 13:1665–1697, May 2012.
- [31] Y. Nesterov. *Introductory Lectures on Convex Optimization*. Kluwer Academic Publishers, New York, 2004.
- [32] Y. Nesterov. Gradient methods for minimizing composite objective function. Technical Report 76, Center for Operations Research and Econometrics (CORE), Catholic University of Louvain (UCL), 2007.
- [33] H. V. Ngai and J. P. Penot. Paraconvex functions and paraconvex sets. *Studia Mathematica*, 184:1–29, 2008.
- [34] G. Raskutti, M. J. Wainwright, and B. Yu. Restricted eigenvalue conditions for correlated Gaussian designs. *Journal of Machine Learning Research*, 11:2241–2259, August 2010.
- [35] G. Raskutti, M. J. Wainwright, and B. Yu. Minimax rates of estimation for high-dimensional linear regression over ℓ_q -balls. *IEEE Trans. Info. Theory*, 57(10):6976–6994, 2011.
- [36] B. Recht. A simpler approach to matrix completion. *Journal of Machine Learning Research*, 12:3413–3430, 2011.
- [37] B. Recht, M. Fazel, and P. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Review*, 52(3):471–501, 2010.
- [38] A. Rohde and A. Tsybakov. Estimation of high-dimensional low-rank matrices. *Annals of Statistics*, 39(2):887–930, 2011.
- [39] M. Rudelson and S. Zhou. Reconstruction from anisotropic random measurements. Technical report, University of Michigan, July 2011.

- [40] N. Srebro, N. Alon, and T. S. Jaakkola. Generalization error bounds for collaborative prediction with low-rank matrices. In *NIPS*, 2005.
- [41] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, 58(1):267–288, 1996.
- [42] J. A. Tropp and A. C. Gilbert. Signal recovery from random measurements via orthogonal matching pursuit. *IEEE Trans. Info. Theory*, 53(12):4655–4666, 2007.
- [43] S. van de Geer and P. Bühlmann. On the conditions used to prove oracle results for the lasso. *Electronic Journal of Statistics*, 3:1360–1392, 2009.
- [44] H. Xu, C. Caramanis, and S. Sanghavi. Robust PCA via outlier pursuit. *IEEE Trans. Info. Theory*, 58(5):3047 –3064, May 2012.
- [45] C. H. Zhang and J. Huang. The sparsity and bias of the lasso selection in high-dimensional linear regression. *Annals of Statistics*, 36(4):1567–1594, 2008.
- [46] P. Zhao, G. Rocha, and B. Yu. Grouped and hierarchical model selection through composite absolute penalties. *Annals of Statistics*, 37(6A):3468–3497, 2009.